RANDOM TREES IN THE BOUNDARY OF OUTER SPACE

ILYA KAPOVICH, JOSEPH MAHER, CATHERINE PFAFF, AND SAMUEL J. TAYLOR

Abstract. We prove that for the harmonic measure associated to a random walk on $\mathrm{Out}(F_r)$ satisfying some mild conditions, a typical tree in the boundary of Outer space is trivalent and nongeometric. This answers a question of M. Bestvina.

1. Introduction

As a means to study the outer automorphism group $\mathrm{Out}(F_r)$, Culler and Vogtmann [CV86] introduced Outer space $\mathrm{CV_r}$ as the deformation space of marked metric F_r-graphs. Outer space is naturally equipped with a boundary $\partial\mathrm{CV_r}$ whose points are represented by actions of F_r on the class of 'very small' $\mathbb{R}$-trees [CL95, BF94]. Since its introduction, $\partial\mathrm{CV_r}$ has attracted much of its own attention and plays a role similar to that of Thurston's boundary of Teichmüller space.

Since a point of $\partial\mathrm{CV_r}$ is the homothety class $[T]$ of an $\mathbb{R}$-tree T, one can study its basic properties as such. For example, each $p \in T$ separates T, and the number of its complementary components is the *valency* of p. We call T *trivalent* if each of its branch-points (i.e. points of valency at least 3) is 3-valent. Similarly, one can also consider the manner in which T arises as an F_r-tree; T is called *geometric* if it is dual to a measured foliation on a 2-complex whose fundamental group is F_r. As a point of reference, *all* of the $\mathbb{R}$-trees that arise in Thurston's boundary of the Teichmüller space are geometric since they are dual to singular measured foliations on the underlying surface. Moreover, in that setting, the valencies of the branch-points correspond to the degrees of the singularities on the surface.

In this paper we develop a complete understanding of these two properties for a "random" tree in $\partial\mathrm{CV_r}$. As a significant point of contrast to the surface case, we find that such a random tree of $\partial\mathrm{CV_r}$ is *not* geometric.

For this, let $(w_n)_{n\geq 1}$ be the random walk on $\mathrm{Out}(F_r)$ determined by a nonelementary measure μ on $\mathrm{Out}(F_r)$. By combining work of Horbez [Hor16] and Namazi–Pettet–Reynolds [NPR14], we recall that the random walk induces a naturally associated *hitting* or *exit* measure ν on $\partial\mathrm{CV_r}$ and that ν is the unique μ-stationary probability measure on $\partial\mathrm{CV_r}$. Moreover, ν gives full measure to the subspace of trees in $\partial\mathrm{CV_r}$ which are free, arational, and uniquely ergodic. We refer the reader to Section 2 for the relevant background. Our main theorem is the following:

Theorem 1.1. *Let $r \geq 3$ and let μ be a nonelementary probability measure on $\mathrm{Out}(F_r)$ with finite support such that the semigroup generated by the support of μ contains φ^{-1} for some principal fully irreducible $\varphi \in \mathrm{Out}(F_r)$.*

Then for ν- almost every $[T] \in \partial\mathrm{CV_r}$, the tree T is trivalent and nongeometric.

This answers a question of Mladen Bestvina, who asked us whether almost every tree in $\partial\mathrm{CV_r}$ is trivalent.

An important component of our argument for Theorem 1.1 is the existence of a *principal* outer automorphism in the semigroup generated by the support of μ. Such outer automorphisms were originally introduced in [AKKP18] and are discussed further in Section 3. Let us remark here that principal outer automorphisms are analogous to pseudo-Anosov mapping classes whose Teichmüller axes live in the top dimensional stratum over Teichmüller space.

As a simple example, we note that the hypotheses of Theorem 1.1 are satisfied when the support of μ is a finite symmetric generating set of $\mathrm{Out}(F_r)$ – see Corollary 7.1 below.

Connections to previous work. In our previous work [KMPT18], we proved that with probability approaching 1 as $n \to \infty$, the random outer automorphism w_n is fully irreducible and its attracting/repelling trees $T_\pm^{w_n}$ are trivalent and nongeometric. However, since such trees form a countable, and hence ν-measure zero, subset of $\partial\mathrm{CV_r}$, this provides no information about a ν-typical tree in $\partial\mathrm{CV_r}$. Indeed, the machinery previously employed, that of ideal Whitehead graphs associated to fully irreducible outer automorphisms, is no longer available in the general setting studied in this paper. Instead, we rely on new results that connect the structure of folding paths to properties of their limiting trees in order to study branching and index properties of the latter.

Our main theorem (Theorem 1.1) in some sense parallels, and is inspired by, the main theorem of [GM17] in the mapping class group setting. There, Gadre–Maher show that with respect to the hitting measure, a typical lamination in Thurston's boundary of Teichmüller space has complementary regions that are triangles and once-punctured disks.

However, our setting differs from theirs in a few key ways. First, their arguments ultimately rely on the openness of the top dimensional stratum in the unit cotangent bundle of Teichmüller space. Of course there is no similar structure for $\mathrm{CV_r}$ and so entirely different techniques must be developed. For this, we introduce the concepts of *eventually legalizing* folding rays (Section 4) and *principal recurrence* (Section 5) which we hope will additionally be useful in future work. Second, as previously mentioned, in the mapping class group setting every limit point of the random walk is geometric (essentially by definition), and so the fact that a typical tree in $\partial\mathrm{CV_r}$ is nongeometric is a truly novel feature of the $\mathrm{Out}(F_r)$-setting. Our argument for this uses the index theory of Gaboriau and Levitt [GL95]. Informally, this states that being nongeometric is equivalent to the failure of a 'Poincaré–Hopf index formula' for branch-points of the tree. Using our specialized folding rays, we show that such a formula typically fails.

Outline of paper. Section 2 provides background on some geometric tools used to study $\mathrm{Out}(F_r)$ and concludes by discussing a few properties of the hitting measure on the boundary of Outer space associated to a random walk on $\mathrm{Out}(F_r)$. In Section 3, we discuss the needed properties of principal outer automorphisms. These are fully irreducible outer automorphisms whose axes in Outer space have particularly rigid and saturated structure. The main result there (Proposition 3.3) says that an arbitrary folding path which closely fellow travels such an axis inherits much of the same structure.

Section 4 presents our main (nonrandom) criteria (Theorem 4.1) ensuring that a folding ray determines a limiting tree that is trivalent and nongeometric. We call such folding paths *eventually legalizing*. Informally, these are folding rays for which every path is, after flowing forward and pulling tight, eventually legal, i.e. no longer folded. If the 'eventually legalizing' condition on the folding ray holds, it allows one to recover the precise structure of the branch-points of the limiting tree T from the graphs along the ray, without losing any directions at the branch-points. A similar issue arose in a recent paper [BHW16], where the authors introduced a "carrying index" of T which sufficed for their purposes but might not detect some directions at branch-points of T.

To establish the eventually legalizing property for a *random* folding ray, we introduce the notion of *principal recurrence* in Section 5. A folding ray is principally recurrent if it fellow travels a translate of a principal axis on arbitrarily long subsegments. The main result (Proposition 5.2) of Section 5 says that random folding rays are principally recurrent.

Finally, in Section 6 we show that a principally recurrent folding path is eventually legalizing (Proposition 6.2). The proof of this fact uses results established in Section 3 and is another instance of a folding path inheriting the structure of a principal axis that it fellow travels. In Section 7 we combine the above results to complete the proof of Theorem 1.1.

Acknowledgments. The first named author was partially supported by NSF grants DMS-1710868 and DMS-1905641. The second named author was supported by Simons Foundation and PSC-CUNY. The third named author acknowledges support from a Queen's University Research Initiation Grant. The last named author thanks Spencer Dowdall for enlightening conversations and is partially supported by NSF grant DMS-1744551. All authors thank M. Bestvina for asking the central question addressed in this paper and acknowledge support from U.S. National Science Foundation grants DMS 1107452, 1107263, 1107367 "RNMS: GEometric structures And Representation varieties" (the GEAR Network).

2. Background

We record here some preliminaries used throughout the paper. Most of this appears in the literature, with exceptions including Proposition 2.1, which builds folding paths to trees in $\partial\mathrm{CV}$, and Corollary 2.3, which establishes that a random tree in $\partial\mathrm{CV}$ is free.

2.1. **Outer space.** We denote by $\widehat{\mathrm{CV}}$ the unprojectivized Outer space for the free group F_r (where $r \geq 2$), and we denote by $\mathrm{CV} = \mathrm{CV_r}$ the corresponding projectivized Outer space. A point in $\widehat{\mathrm{CV}}$ is represented (up to some natural equivalence) by a marked metric graph structure on a finite connected graph G where each vertex of

G has degree ≥ 3, the *metric* assigns each edge of G a strictly positive length, and the *marking* identifies $\pi_1(G)$ with F_r. We can also think of this point of $\widehat{\mathrm{CV}}$ as the minimal free discrete isometric action of F_r on the $\mathbb{R}$-tree $T = \widetilde{G}$ with the lifted metric. We denote by $\mathrm{vol}(G) = \mathrm{vol}(T)$ the sum of the lengths of the edges of G. The space $\mathrm{CV} \subseteq \widehat{\mathrm{CV}}$ consists of points $G \in \widehat{\mathrm{CV}}$ with $\mathrm{vol}(G) = 1$.

There is a natural closure $\overline{\widehat{\mathrm{CV}}}$ of $\widehat{\mathrm{CV}}$ with respect to the length function topology, and $\overline{\widehat{\mathrm{CV}}}$ is known to consist of precisely the *very small* nontrivial minimal isometric actions by F_r on $\mathbb{R}$-trees. The projectivization of $\overline{\widehat{\mathrm{CV}}}$ with respect to the natural multiplication action of $\mathbb{R}_{>0}$ is denoted $\overline{\mathrm{CV}}$; it is known that $\overline{\mathrm{CV}}$ is compact. For every $T \in \widehat{\mathrm{CV}}$ the projective class $[T]$ is canonically identified with $T/\mathrm{vol}(T) \in \mathrm{CV}$, and thus we can think of CV as the projectivization of $\widehat{\mathrm{CV}}$, and so as a subset of $\overline{\mathrm{CV}}$. We denote $\partial\mathrm{CV} = \overline{\mathrm{CV}} - \mathrm{CV}$. For additional background on Outer space, its topology, and its boundary see [CV86, CL95, BF94, Pau89].

For $G_1, G_2 \in \widehat{\mathrm{CV}}$, we denote by $\Lambda(G_1, G_2)$ the infimum of the Lipschitz constants of the continuous maps $f \colon G_1 \to G_2$ preserving the marking, i.e. "change of marking" maps. It is known that for $G_1, G_2 \in \mathrm{CV}$ we have $\Lambda(G_1, G_2) \geq 1$, and that $G_1 = G_2$ in CV if and only if $\Lambda(G_1, G_2) = 1$. For $G_1, G_2 \in \mathrm{CV}$ we denote $d_{\mathrm{CV}}(G_1, G_2) = \log \Lambda(G_1, G_2)$ and refer to d_{CV} as the *asymmetric Lipschitz metric* on CV. For more on this metric, see [FM11, AK11, BF14]. As is common, we let d_{sym} denote the *symmetric Lipschitz metric*: $d_{\mathrm{sym}}(G_1, G_2) = d_{\mathrm{CV}}(G_1, G_2) + d_{\mathrm{CV}}(G_2, G_1)$.

For an interval $J \subseteq \mathbb{R}$, a map $\gamma \colon J \to \mathrm{CV}$ is called a *geodesic* in CV if $d_{\mathrm{CV}}(\gamma(t), \gamma(t')) = t' - t$ for all $t, t' \in J$ with $t \leq t'$. A *geodesic ray* in CV is a geodesic $\gamma \colon [0, \infty) \to \mathrm{CV}$. We emphasize that the term geodesic always refers to the asymmetric Lipschitz metric.

2.2. **Laminations and arational trees.** We refer the reader to [CHL08a, CHL08b, Rey12, BR15, BF14] for detailed background on algebraic laminations on F_r, arational trees, and the free factor complex. We only recall a few basic facts here. For a free group F_r (with $r \geq 2$) let $\partial^2 F_r = \{(z_1, z_2) \in \partial F_r \times \partial F_r | z_1 \neq z_2\}$. The set $\partial^2 F_r$ is equipped with the subspace topology from $\partial F_r \times \partial F_r$ and with the diagonal translation action of F_r. An *algebraic lamination* on F_r is a subset $L \subseteq \partial^2 F_r$ which is closed, F_r-invariant, and flip-invariant (for the "flip" map $\partial^2 F_r \to \partial^2 F_r$ defined by $(z_1, z_2) \mapsto (z_2, z_1)$). For an algebraic lamination L on F_r a pair $(z_1, z_2) \in L$ is called a *leaf* of L. For a lamination L on F_r, a leaf $(z_1, z_2) \in L$, and a nontrivial finitely generated subgroup $H \leq F_r$ we say that (z_1, z_2) is *carried* by H if both z_1 and z_2 are contained in ∂H.

For any tree $T \in \overline{\widehat{\mathrm{CV}}}$ there is an associated *dual lamination* or *zero lamination* $L(T) \subseteq \partial^2 F_r$ on F_r which depends only on the projective class $[T] \in \overline{\mathrm{CV}}$. The dual lamination encodes, in a systematic way, the information about sequences of elements of F_r with arbitrarily small translation length in T. We refer the reader to [CHL08b] for the precise technical definition of $L(T)$. For our purposes the key relevant facts are that for $T \in \overline{\widehat{\mathrm{CV}}}$ we have $L(T) = \varnothing$ if and only if $T \in \widehat{\mathrm{CV}}$, and that whenever $T, T' \in \overline{\widehat{\mathrm{CV}}}$ are such that $||u||_T \leq ||u||_{T'}$ for every $u \in F_r$ then $L(T') \subseteq L(T)$. A tree $T \in \overline{\mathrm{CV}}$ is called *arational* if $T \notin \widehat{\mathrm{CV}}$ and if no leaf of $L(T)$ is carried by a proper free factor of F_r [Rey12]. In this case the projectivized tree $[T] \in \partial\mathrm{CV}$ is also called *arational.* Note that the property of being arational depends only on the dual lamination of the tree.

For $r \geq 3$, the *free factor graph* $\mathcal{FF}$ is a simple graph where the vertex set is the set of F_r-conjugacy classes of proper free factors of F_r. Two distinct vertices of $\mathcal{FF}$ are adjacent in $\mathcal{FF}$ if and only if they can be represented as conjugacy classes $[A], [B]$ of proper free factors A, B of F_r such that $A \leq B$ or $B \leq A$. The graph $\mathcal{FF}$ is endowed with the simplicial metric where every edge has length 1, and with the natural left action of $\mathrm{Out}(F_r)$ by simplicial automorphisms (and hence by isometries), where for a vertex $[A]$ of $\mathcal{FF}$ and an element $\phi \in \mathrm{Out}(F_r)$ we have $\phi \cdot [A] = [\phi(A)]$.

It is known, by a result of Bestvina and Feighn [BF14], that for $r \geq 2$ the free factor graph $\mathcal{FF}$ is Gromov-hyperbolic, and that for $\phi \in \mathrm{Out}(F_r)$ the element ϕ acts as a loxodromic isometry if and only if ϕ is fully irreducible. There is a natural coarsely defined and coarsely $\mathrm{Out}(F_r)$-equivariant "projection" $\pi\colon \mathrm{CV} \to \mathcal{FF}$ where $G_0 \in \mathrm{CV}$ is mapped to the free factor $[A]$ represented by any proper connected non-contractible subgraph of G_0. It is also known [BR15] (see also [Ham12]) that the hyperbolic boundary $\partial\mathcal{FF}$ can be identified with the set of equivalence classes $[[T]]$ of arational trees $T \in \overline{\mathrm{CV}}$, where two such trees T, T' are considered equivalent whenever $L(T) = L(T')$.

Finally, let $\mathcal{UE}$ be the subspace of $\partial\mathrm{CV}$ consisting of arational trees having a unique length measure, up to scale. More precisely, $[T] \in \mathcal{UE}$ if and only if T is arational and $[T] = [T']$ whenever $L(T) = L(T')$. Such trees are sometimes called *uniquely ergodic.*

2.3. **Branch-points and the geometric index of a tree.** For an $\mathbb{R}$-tree T and a point $p \in T$, a *direction* at p in T is a connected component of $T \setminus \{p\}$. The number of directions at p in T is denoted $\mathrm{val}_T(p)$ and called the *valency* (or *degree*) of p in T. We think of $\mathrm{val}_T(p)$ as an element of $\{\infty\} \cup \{n \in \mathbb{Z} | n \geq 0\}$. A point $p \in T$ is a *branch-point* of T if $\mathrm{val}_T(p) \geq 3$.

Let $T \in \overline{\mathrm{CV}}$. In [GL95] Gaboriau and Levitt proved that T has only finitely many F_r-orbits of branch-points and only finitely many F_r-orbits of directions at branch-points. They also showed that if $T \in \overline{\mathrm{CV}}$ is a free F_r-tree then for every branch-point $p \in T$ one has $\mathrm{val}_T(p) < \infty$. For such a free F_r-tree T, if $p_1, \dots, p_m \in T$ are representatives of all the distinct F_r-orbits of branch-points, [GL95] defined the *geometric index* $\mathrm{ind}_{geom}(T)$ as

$$\mathrm{ind}_{geom}(T) = \sum_{i=1}^{m} [\mathrm{val}_T(p_i) - 2].$$

The unordered list $\mathrm{val}_T(p_1), \dots, \mathrm{val}_T(p_m)$ is the *index list* for T.

Gaboriau and Levitt further defined $\mathrm{ind}_{geom}(T)$ for an arbitrary (not necessarily free) tree $T \in \overline{\mathrm{CV}}$ and proved that one always has $\mathrm{ind}_{geom}(T) \leq 2r - 2$. The equality $\mathrm{ind}_{geom}(T) = 2r - 2$ holds if and only if the tree T is *geometric*, i.e. arises as the dual tree of a measured foliation of some finite 2-complex with fundamental group F_r. We say that T is *nongeometric* if $\mathrm{ind}_{geom}(T) < 2r-2$. We refer the reader to the paper [CH12] for more detailed background on this topic.

2.4. **Folding lines and limiting trees.** We next turn to folding paths in CV and in $\widehat{\mathrm{CV}}$. In the case of folding paths between simplicial trees, we closely follow [BF14, Section 2], where we refer the reader for additional details. Since we will be particularly interested in folding rays to points in $\partial\mathrm{CV}$, we pay special attention to this case in Proposition 2.1.

Following [HM11, MP16], we define a *folding path* in $\widehat{\mathrm{CV}}$ as a proper continuous injective map $\gamma\colon I \to \widehat{\mathrm{CV}}$ (where $I \subseteq \mathbb{R}$ is an interval), with $\gamma(t) = G_t \in \widehat{\mathrm{CV}}$ for all $t \in I$, together with a family of continuous *folding maps* $g_{t,t'}\colon G_t \to G_{t'}$, where $t, t' \in I$ with $t \le t'$, satisfying the following properties: Each map $g_{t,t'}\colon G_t \to G_{t'}$ is locally injective on edges of G_t, and we have $g_{t,t} = Id_{G_t}$ for each $t \in I$. In addition, whenever $t \le t' \le t''$ for $t, t', t'' \in I$, we have $g_{t,t''} = g_{t',t''} \circ g_{t,t'}$. We will often denote such a folding path as just $(G_t)_{t\in I}$ and suppress explicit mention of the maps $g_{t,t'}$. A folding path is a *folding line* if $I = \mathbb{R}$ and a *folding ray* if $I = [t_0, \infty)$ for some $t_0 \in \mathbb{R}$.

For the most part, in this paper we will concentrate on special "greedy" types of folding paths. We next turn to their description and refer the reader to [BF14, FM11] for more details.

For a point $G \in \widehat{\mathrm{CV}}$, a *gate structure* $\mathcal{T}$ on G is a partition, for every vertex v of G, of the set of oriented edges originating at v into nonempty subsets called *gates*. A turn $\{e_1, e_2\}$ at v is called *legal* with respect to $\mathcal{T}$ if e_1, e_2 belong to different gates, and is called *illegal* otherwise. In this setting the gate structure and the notions of legal and illegal turns naturally extend, via lifting, to $T = \widetilde{G}$. An edge-path (or a circuit) in G is called *legal* with respect to $\mathcal{T}$ if for every 2-edge subpath ee' of this path, the turn $\{e^{-1}, e'\}$ is legal. A *train track structure* on G is a gate structure $\mathcal{T}$ on G such that at each vertex of G there are at least 2 gates.

For trees $T_0 \in \widehat{\mathrm{CV}}$, $T \in \overline{\widehat{\mathrm{CV}}}$, an F_r-equivariant map $f\colon T_0 \to T$ is called a *morphism* if for each edge $e = [x, y]$ of T the map f sends e isometrically to $[f(x), f(y)]_T$ (so that, in particular, $f(x) \ne f(y)$). Note that a morphism is, by definition, a 1-Lipschitz map. A morphism $f\colon T_0 \to T$ defines a *pullback* gate structure $\mathcal{T}_f$ on T_0 where a turn $\{e_1, e_2\}$ at a vertex x of T_0 is legal if and only if the restriction of the map f to the path $e_1^{-1}e_2$ is injective. A morphism $f\colon T_0 \to T$ is *optimal* if the pullback gate structure $\mathcal{T}_f$ is a train track structure on T_0.

Suppose $T_0 = \widetilde{G}_0 \in \widehat{\mathrm{CV}}$, $T \in \overline{\widehat{\mathrm{CV}}}$, and $f\colon T_0 \to T$ is an optimal morphism. Then f canonically determines in $\widehat{\mathrm{CV}}$ a *greedy isometric folding path defined by* f, denoted $(\widehat{G}_s)_{s\in J}$, with $J \subseteq [0, \infty)$ an interval starting at 0, with $G_0 = \widehat{G}_0$, and with the following properties and additional structure. For every $s, s' \in J$ with $s \le s'$ we have a 1-Lipschitz map $\widehat{g}_{s,s'}\colon \widehat{G}_s \to \widehat{G}_{s'}$ that lifts to an optimal morphism $f_{s,s'}\colon T_s \to T_{s'}$, where $T_s = \widetilde{\widehat{G}}_s$ and $T_{s'} = \widetilde{\widehat{G}}_{s'}$. For each $s \in J$ we also have an optimal morphism $f_s\colon T_s \to T$, where $f_0 = f$. These morphisms are compatible, in the sense that for every $s, s' \in J$ with $s \le s'$ we have $f_{s'} \circ f_{s,s'} = f_s$. For each $s \in J$ we equip T_s with the pullback gate structure $\mathcal{T}_s$ induced by $f_s\colon T_s \to T$. (In what follows, we will refer to both sets of maps $\widehat{g}_{s,s'}$ and $f_{s,s'}$ as *folding maps*.) The "greedy" property of this folding line means that for each $s \in J$, which is not the right-end point of J, there exists an $\epsilon > 0$ such that $[s, s+\epsilon) \subseteq J$ and such that for each $s' \in (s, s+\epsilon)$ the map $f_{s,s'}\colon T_s \to T'_s$ is obtained by equivariantly, at each vertex x of T_s and for each gate (with respect to $\mathcal{T}_s$) at x, folding together into a single segment the initial segments of length $s'-s$ of all the edges in that gate. The interval J starting at 0 is chosen to be maximal possible subject to $(\widehat{G}_s)_{s\in J}$ satisfying all these properties.

For several constructions of greedy folding lines and additional properties, see [BF14, Section 2]. We remark on a few relevant properties here. The function $\mathrm{vol}(T_s)$ is strictly monotone decreasing on J. Moreover, the fact that $f\colon T_0 \to T$ is an optimal morphism implies that for each $s \in J$ the pullback gate structure $\mathcal{T}_s$ on $\widehat{G}_s$ is a train track structure. The path $(\widehat{G}_s)_{s\in J}$, with the maps $\widehat{g}_{s,s'}$, is a folding path in CV in the more general sense described in Subsection 2.4. Also, in this setting, for any $s_1 \le s_2$ in J the path $(\widehat{G}_s)_{s\in[s_1,s_2]}$ is (up to shifting the parameter by s_1) exactly the greedy isometric folding path defined by $f_{s_1,s_2}\colon T_{s_1} \to T_{s_2}$.

It is known that if $f\colon T_0 \to T$ is an optimal morphism, then the path $(\widehat{G}_s)_{s\in J}$ projects to a reparameterized geodesic in CV [FM11, AK11]. In this case for $s, s' \in J$ with $s \le s'$ we have $\widehat{G}_s/\mathrm{vol}(\widehat{G}_s), \widehat{G}_{s'}/\mathrm{vol}(\widehat{G}_{s'}) \in \mathrm{CV}$ and

$$d_{\mathrm{CV}}\left(\frac{\widehat{G}_s}{\mathrm{vol}(\widehat{G}_s)}, \frac{\widehat{G}_{s'}}{\mathrm{vol}(\widehat{G}_{s'})}\right) = \log \frac{\mathrm{vol}(\widehat{G}_s)}{\mathrm{vol}(\widehat{G}_{s'})}.$$

In particular, if $G_0 \in \mathrm{CV}$ has volume 1, then in this setting

$$d_{\mathrm{CV}}\left(G_0, \frac{\widehat{G}_s}{\mathrm{vol}(\widehat{G}_s)}\right) = \log \frac{1}{\mathrm{vol}(\widehat{G}_s)} = -\log \mathrm{vol}(\widehat{G}_s).$$

Since $\mathrm{vol}(\widehat{G}_s)$ is a strictly decreasing function on J, there exists a unique monotone increasing reparameterization $\alpha(t)$ of J with $\alpha(0) = 0$, $\alpha\colon J' \to J$, such that $\mathrm{vol}(\widehat{G}_{\alpha(t)}) = e^{-t}$ for all $t \in J'$. We denote $G_t = \widehat{G}_{\alpha(t)}/\mathrm{vol}(\widehat{G}_{\alpha(t)})$ for all $t \in J'$. Note that as topological spaces we have $G_t = \widehat{G}_{\alpha(t)}$, and the only difference between G_t and $\widehat{G}_{\alpha(t)}$ is in their metric graph structures. For all $t \le t'$ in J' we also set $g_{t,t'} = \widehat{g}_{\alpha(t),\alpha(t')}$. Then $(G_t)_{t\in J'}$, with the maps $g_{t,t'}$, is a folding path in CV in the general sense described above.

This reparameterization gives us a path $(G_t)_{t\in J'}$ in CV starting at G_0 which is a geodesic in CV. If $G_0 \in \mathrm{CV}$, $T_0 = \widetilde{G}_0$, $T \in \overline{\widehat{\mathrm{CV}}}$, and $f\colon T_0 \to T$ is an optimal morphism, we refer to $(G_t)_{t\in J'}$ as the *greedy geodesic folding path defined by* f.

If $T \in \mathrm{CV}$, then in the above setting a greedy geodesic folding path defined by f always reaches T in some finite time, and $J' = [0, d_{\mathrm{CV}}(T_0, T)]$. If $[T] \in \partial\mathrm{CV}$, then it is possible that J' is a finite interval (this can happen if the geodesic folding path exits CV after a finite distance), and even in the case where $J' = [0,\infty)$ we are not necessarily guaranteed that $\lim_{t\to\infty} G_t = [T]$ in $\overline{\mathrm{CV}}$. Nevertheless, for reasonably nice $T \in \partial\mathrm{CV}$ one can rule out such unexpected behavior.

Proposition 2.1. *Let* $[T] \in \partial\mathrm{CV}$ *be such that* T *is a free* F_r*-tree. Then:*

(1) For each r-rose in CV *there exists a metric structure $G_0 \in \mathrm{CV}$ on this rose and an optimal morphism $f\colon \widetilde{G}_0 = T_0 \to cT$ for some $c > 0$.*

(2) Let $T_0 = \widetilde{G}_0 \in \mathrm{CV}$, let $f\colon T_0 \to T$ be an optimal morphism, and let $(\widehat{G}_s)_{s\in J}$ and $(G_t)_{t\in J'}$ be the greedy isometric folding path and the greedy geodesic folding path determined by f. Then:

(a) There exists a limit $\lim_{s\to\infty} \widehat{G}_s = T'$ in $\overline{\widehat{\mathrm{CV}}}$, and, moreover, T' is again a free F_r-tree and $[T'] \in \partial\mathrm{CV}$. Moreover, in this case $L(T') \subseteq L(T)$.

(b) If, in addition, T is arational, then $L(T) = L(T')$ and $J' = [0,\infty)$, so that

$$\lim_{s\to\infty} d_{\mathrm{CV}}\left(G_0, \frac{\widehat{G}_s}{\mathrm{vol}(\widehat{G}_s)}\right) = \infty$$

(c) If T is arational and uniquely ergodic, then $T' = T$ in $\overline{\widehat{\mathrm{CV}}}$, and hence

$$\lim_{s\to\infty} \widehat{G}_s = T$$

in $\overline{\mathrm{CV}}$.

Proof. (1) Let $\Gamma_0 \in \mathrm{CV}$ be an r-rose corresponding to a free basis $a_1, \dots, a_r$ of F_r. By assumption F_r acts freely on T, so that a_1 is a loxodromic isometry of T with translation length $||a_1||_T > 0$.

Let $x_0 \in \widetilde{\Gamma}_0$ be a lift of the vertex v_0 of Γ_0. Let $L_{a_1} \subseteq T$ be the axis of a_1 in T, and pick a point $p \in L_{a_1}$. Thus $a_1 p \in L_{a_1}$ and $d_T(p, a_1 p) = ||a_1||_T > 0$. By replacing T by cT for an appropriate $c > 0$ we can assume that $\sum_{i=1}^r d_T(p, a_i p) = 1$.

Note that since T is a free F_r-tree, we have $a_i p \neq p$ for $i = 1, \dots, r$. We give each edge a_i of Γ_0 the length $d_T(p, a_i p) > 0$, which defines a new volume-1 metric structure G_0 on Γ_0, and a point $T_0 = \widetilde{G}_0 \in \mathrm{CV}$. For $i = 1, \dots, r$ denote by x_i the vertex of T which is the terminal endpoint of the lift e_i of the petal a_i of Γ_0 starting at x_0. We construct an F_r-equivariant morphism $f\colon T_0 \to T$ by setting $f(x_0) = p$, setting $f(x_i) = a_i p$ for $i = 1, \dots, r$, mapping each e_i isometrically to the segment $[p, a_i p]_T$, and then extending f by equivariance. By construction $f\colon T_0 \to T$ is a morphism. Moreover, the fact that $p \in L_{a_1}$ implies that x_0 (and hence every other vertex of T_0) has at least 2 gates for the pullback gate structure $\mathcal{T}_f$. Thus f is an optimal morphism, as required.

(2)

(a) Since $f\colon T_0 \to T$ is an optimal morphism, hence each vertex for the pullback legal structure $\mathcal{T}_0$ on T_0 has at least 2 gates at each vertex, there exists a nontrivial $\mathcal{T}_0$-legal circuit γ in G_0 representing the conjugacy class of some $1 \neq w \in F_r$. The fact that $(\widehat{G}_s)_{s\in J'}$ is the greedy isometric folding line determined by f and starting at $G_0 = \widehat{G}_0$ implies that for each $s \in J$ the circuit $f_{0,s}(\gamma)$ is legal in $\widehat{G}_s$ for the train track structure $\mathcal{T}_s$ induced by $f_s\colon T_s \to T$. Denote $M = \sup\{s \mid s \in J\}$. Thus $0 < M < \mathrm{vol}(G_0) < \infty$.

The fact that for any $s \leq s'$ in J the folding map $f_{s,s'}\colon T_s \to T_{s'}$ is 1-Lipschitz implies that for each $u \in F_r$ we have $||u||_{T_s} \geq ||u||_{T_{s'}}$. Thus for each $u \in F_r$ the function $||u||_{T_s}$ is monotone non-increasing on J and there is a finite limit $\lim_{s\to M^-} ||u||_{T_s}$.

Moreover, for our legal loop γ representing $1 \neq w \in F_r$ we have $||w||_{T_s} = ||w||_{T_0} > 0$, and so the limit $\lim_{s\to M-} ||w||_{T_s} = ||w||_{T_0} > 0$. Therefore there exists a nontrivial tree $\lim_{s\to\infty} T_s = T'$ in $\overline{\widehat{\mathrm{CV}}}$. Since there are 1-Lipschitz maps $f_s \colon T_s \to T$, we have $||u||_{T_s} \geq ||u||_T$ for every $u \in F_r$ and every $s \in J$. Therefore, for the limiting length function $||.||_{T'}$, we also have $||u||_{T'} \geq ||u||_T$ for all $u \in F_r$. Recall that T is a free F_r-tree. Therefore for every $1 \neq u \in F_r$ we have $||u||_{T'} \geq ||u||_T > 0$, so that T' is also a free F_r-tree.

We claim that $[T'] \in \partial\mathrm{CV}$. Suppose not. Then $T' \in \widehat{\mathrm{CV}}$ and $\sup_J s = M \in J$ and $T' = T_M$. The assumption that $T \in \partial\mathrm{CV}$ then implies that the map $f_M \colon T_M \to T$ is not locally injective, and therefore for the gate structure $\mathcal{T}_M$ on T_M there exists a gate at some vertex with at least two distinct edges in that gate. This means that the isometric folding path $(\widehat{G}_s)_{s\in J}$ can be continued past $s = M$ for some positive time $[M, M+\epsilon)$, contradicting the fact that $M = \sup_J s$. The condition $||u||_{T'} \geq ||u||_T$ for all $u \in F_r$ also implies that $L(T') \subseteq L(T)$. This completes the proof of (2)(a).

(b) Suppose now that, in addition, T is both free and arational. By part (a) above we know that $[T'] \in \partial\mathrm{CV}$ and therefore $L(T') \neq \varnothing$. Now [BR15, Proposition 4.2(i)] implies that the "derived lamination" $L'(T) \subseteq L(T)$ is the unique minimal sublamination in $L(T)$. Since $L'(T)$ is minimal, we have $L'(T) = L''(T) = L'''(T)$. Since $L(T') \subseteq L(T)$, and since $L(T')$ is a nonempty lamination, it follows that $L'(T) \subseteq L(T')$. Thus $L'''(T) \subseteq L(T')$. Since T is arational, [BR15, Corollary 4.3] implies that $L(T') = L(T)$, and that T' is also arational.

Then the greedy geodesic folding path $(G_t)_{t\in J'}$ projects to a reparameterized quasi-geodesic in the free factor complex $\mathcal{FF}$ [BF14, Corollary 6.5] which converges to a point of the hyperbolic boundary $\partial\mathcal{FF}$ represented by T [BR15, Proposition 8.3]. Since the projection map $\pi\colon \mathrm{CV} \to \mathcal{FF}$ is coarsely Lipschitz, it follows that $J' = [0, \infty)$. Indeed, otherwise J' is a finite interval and π would map the folding line $(G_t)_{t\in J'}$ to a set of bounded diameter in $\mathcal{FF}$, which cannot limit to a point of $\partial\mathcal{FF}$. Thus indeed $J' = [0, \infty)$ and $\lim_{s\to\infty} d_{\mathrm{CV}}(G_0, \widehat{G}_s/\mathrm{vol}(\widehat{G}_s)) = \infty$. Part (2)(b) is verified.

(c) Suppose now that T is free arational and uniquely ergodic. By part (b) we know that $L(T) = L(T')$ and T' is arational. Then, by definition of unique ergodicity, we have $[T] = [T']$ in $\partial\mathrm{CV}$. Thus $T' = bT$ for some $b > 0$. Note that for our legal circuit γ representing w in G_0 we have $||w||_T = ||w||_{T_0} = ||w||_{T'} > 0$ and therefore $b = 1$. Thus $T = T'$ in $\overline{\widehat{\mathrm{CV}}}$, as required. □

We conclude this subsection by setting a few conventions to simplify terminology.

Convention 2.2. From now on, by a *geodesic folding ray* in CV we mean a folding ray $(G_t)_{t\in[t_0,\infty)}$ in CV which, up to a shift of the parameter by t_0, is a greedy geodesic folding path in CV with $J' = [0, \infty)$. Also, by a *geodesic folding line* in CV we mean a folding line $(G_t)_{t\in\mathbb{R}}$ in CV such that for every $t_0 \in R$ the path $(G_t)_{t\in[t_0,\infty)}$ is a geodesic folding ray in CV.

We will often abbreviate the notation for geodesic folding rays and geodesic folding lines in CV to just (G_t). Moreover, if a geodesic folding line in CV is φ-periodic for some fully irreducible $\varphi \in \mathrm{Out}(F_r)$, we usually denote such a line by $A(t)$.

2.5. **Random walks and Outer space.** The general notion of a nonelementary probability measure on a group acting isometrically on a Gromov-hyperbolic metric space is discussed in more detail in Section 5 below. Considering the case of the action of $G = \mathrm{Out}(F_r)$ on the free factor graph $\mathcal{FF}$, a probability measure μ on $\mathrm{Out}(F_r)$ is *nonelementary* if the subsemigroup $\langle \mathrm{Supp}(\mu)\rangle_+$ of $\mathrm{Out}(F_r)$ generated by the support of μ contains some two independent fully irreducible elements ψ_1, ψ_2. Here *independent* means that the attracting and repelling fixed points of ψ_1, ψ_2 in $\partial\mathcal{FF}$ are four distinct points. By [BFH97, Proposition 2.16, Theorem 4.1], fully irreducibles $\psi_1, \psi_2 \in \mathrm{Out}(F_r)$ are independent if and only if $\langle \psi_1, \psi_2\rangle \leq \mathrm{Out}(F_r)$ is not virtually cyclic, and also if and only if $\langle \psi_1\rangle \cap \langle \psi_2\rangle = \{1\}$.

Recall that $\mathcal{UE} \subset \partial\mathrm{CV}$ is the subspace of uniquely ergodic trees.

The following is Theorem 7.21 of Namazi–Pettet–Reynolds [NPR14]; see also Dahmani–Horbez [DH17, Theorem 5.10] and Horbez [Hor17, Proposition 4.4].

Theorem 2.1 (Hitting measure on $\partial\mathrm{CV}$)**.** *Let μ be a nonelementary probability measure on $\mathrm{Out}(F_r)$ with finite first moment with respect to d_{CV}. Then for almost every sample path $\omega = (\omega_n)_{n\geq 0}$ of the random walk on $(\mathrm{Out}(F_r), \mu)$ and any $y_0 \in \mathrm{CV}$, the sequence $(\omega_n y_0)_{n\geq 0}$ converges to a point $\mathrm{bnd}(\omega) \in \mathcal{UE}$. The hitting measure ν defined by setting*

$$\nu(S) = \mathbb{P}(\mathrm{bnd}(\omega) \in S),$$

for all measurable subsets $S \subset \partial\mathrm{CV}$ is nonatomic, and it is the unique μ-stationary measure on $\partial\mathrm{CV}$.

In fact, it is not hard to see that ν–almost every $T \in \partial\mathrm{CV}$ is also free. Since we will need this fact, we record it here.

Corollary 2.3. *Suppose in addition to the hypotheses of Theorem 2.1 that the semigroup generated by the support of μ contains a nongeometric fully irreducible automorphism. Let ν be the associated hitting measure on $\partial\mathrm{CV}$ as obtained in Theorem 2.1. Then a ν-typical tree T in $\partial\mathrm{CV}$ is free.*

Proof. The hypotheses imply that μ is nonelementary with respect to the action on the co-surface graph (See [TT16, Section 2.3]). By Maher–Tiozzo [MT14, Theorem 1.1], this means that almost every sample path converges to a point in the boundary of the co-surface graph. By work of Dowdall–Taylor [DT17] the boundary of the co-surface graph is the subspace of $\partial\mathrm{CV}$ consisting of free and arational trees (after identifying trees with the same dual lamination, as in the identification of $\partial\mathcal{FF}$).

Now for a typical sample path ω, $(\omega_n y_0)_{n\geq 0}$ converges to a point $\mathrm{bnd}(\omega) \in \mathcal{UE}$ by Theorem 2.1. Since such a path typically projects to a path in the co-surface graph converging to a boundary point represented by a free tree, we see that $\mathrm{bnd}(\omega)$ is also free. $\square$

The additional assumption in Corollary 2.3 on the semigroup generated by the support of μ is necessary. Without it, the entire random walk could, for example, be contained in some mapping class subgroup of $\mathrm{Out}(F_r)$ in which case almost every limiting tree has nontrivial point stabilizers.

3. Principal outer automorphisms and fellow traveling folding paths

We now turn to discussing the particular type of outer automorphism, called a *principal* outer automorphism, that will act as the 'seed' of our construction. The main result of this section (Proposition 3.3) proves a strong rigidity property for folding paths that fellow travel the axis of a principal outer automorphism.

The original definition of a principal outer automorphism $\varphi \in \mathrm{Out}(F_r)$ is given in terms of its ideal Whitehead graph [HM11] and the reader can find a complete definition in those terms in [AKKP18] or [KMPT18]. Rather than recall the original definition here, we collect the essential properties that we will need and give an alternative characterization.

Recall that a fully irreducible $\varphi \in \mathrm{Out}(F_r)$ is called *ageometric* if the attracting tree $T_+ = T_+^\varphi \in \partial\mathrm{CV}$ is nongeometric, i.e. $\mathrm{ind}_{geom} T_+^\varphi < 2r-2$. For an ageometric fully irreducible $\varphi \in \mathrm{Out}(F_r)$ the action of F_r on T_+^φ is free and has dense F_r-orbits. For $r \geq 3$, a fully irreducible $\varphi \in \mathrm{Out}(F_r)$ is *principal* if φ is ageometric with $\mathrm{ind}_{geom} T_+^\phi = 2r-3$, if every branch-point $p \in T_+^\varphi$ has $\mathrm{val}_{T_+}(p) = 3$, and if every nondegenerate turn at p in T_+^φ is "taken" by the expanding lamination Λ_ϕ of φ. For those readers unacquainted with this terminology, this notion essentially amounts to the fact that among all fully irreducible outer automorphisms, principal outer automorphisms are characterized as those which satisfy conditions $(2)-(4)$ in Lemma 3.1. We remark that principal outer automorphisms exist in $\mathrm{Out}(F_r)$ for each $r \geq 3$ [AKKP18, Example 6.1].

As a fully irreducible outer automorphism, a principal $\varphi \in \mathrm{Out}(F_r)$ has a periodic folding line A in CV, which we write as $A(t)$ rather than (A_t) as done in Section 2.4. Here, A is periodic in the sense that there is a $\lambda > 1$ so that $\varphi^{-1}A(t) = A(t)\cdot\varphi = A(t+\ln\lambda)$ for all $t \in \mathbb{R}$ [1]. Note that $\ln\lambda > 0$ is the translation length of φ in CV. We refer to A as an *axis* for φ.

Next we collect properties of the pair φ, A. Most of these are easily located in the literature.

Lemma 3.1. *Suppose that $\varphi \in \mathrm{Out}(F_r)$ is principal and that A is an axis for φ.*

1. *The folding line A is the* lone axis *for φ. This means that it is the unique (up to reparameterization) folding line with the property that $\lim_{t\to-\infty} A(t) = [T_-]$ and $\lim_{t\to\infty} A(t) = [T_+]$, where $[T_-],[T_+] \in \partial\mathrm{CV}$ are the repelling/ attracting trees for φ.*
2. *For all but a discrete collection of times, $A(t)$ is contained in the interior of a maximal simplex (i.e. it is trivalent). Moreover, when $A(t)$ is not trivalent, it has a unique vertex of degree 4.*
3. *For all $t \in \mathbb{R}$, $A(t)$ has exactly one illegal turn. Hence, A is a greedy folding line in the sense defined in Section 2.4.*

[1]Note that it is φ^{-1} that translates along the forward 'folding' direction of A for the left action on CV.

(4) *For all $t \in \mathbb{R}$ for which $A(t)$ is trivalent, every legal turn of $A(t)$ is taken (i.e. it is a turn traversed by the image of the interior of an edge of $A(s)$ under the folding map $A(s) \to A(t)$ for some $s < t$).*

Proof. Since φ is a principal outer automorphism, by definition, its ideal Whitehead graph $IW(\varphi)$ is the disjoint union of $2r-3$ triangles. Thus, (1) is a direct consequence of [MP16, Theorem 4.7] and the [HM11] definition of an axis bundle.

Similarly, item (2) follows immediately from Lemma 5.1 and Remark 3.11 in [AKKP18], and item (3) is explained in [KMPT18, Remark 5.4] using the fact that A is a lone axis for φ (as in item (1)).

To prove item (4), recall that in the language of Section 2.4, $A(t)$ (for t greater than any fixed $t_0 \in \mathbb{R}$) is a greedy geodesic folding path guided by some optimal morphism $f\colon \widetilde{A(t_0)} \to T_+$, where T_+ is the attracting tree for φ (as in item (1)). We suppose that $A(t_0)$ is trivalent and let v_0 be its unique vertex with an illegal turn (using item (3)). For any other vertex v of $A(t_0)$ and any lift $\widetilde{v}$ to $\widetilde{A(t_0)}$, f maps $\widetilde{v}$ to a (necessarily valence 3) branch-point of T_+. From the property that $\mathrm{ind}_{geom} T_+ = 2r-3$ we note that f induces a bijection between the set of vertices of $A(t_0)$ other than v_0 and the set of orbits of branch-points of T_+. The condition that all nondegenerate turns at $f(\widetilde{v})$ are 'taken' by the stable lamination means here that for each such turn there is an edge $\widetilde{e}$ of $\widetilde{A(t_0)}$ whose interior maps over this turn under f. In terms of the greedy geodesic folding line A, this translates to the statement that for some sufficiently large integer n, the folding map $A(t_0) \to A(t_0 + n\ln\lambda) = A(t_0)\cdot\varphi^n$ has the property that the image of each vertex $v \neq v_0$, which is itself a trivalent vertex with all legal turns, has each of its turns taken by some edge of $A(t_0)$.

Since t_0 was an arbitrary time for which $A(t_0)$ is trivalent, using periodicity of the folding line A we see that it only remains to show that the two legal turns of v_0 are taken by edges of $A(s)$ under the folding map $A(s) \to A(t_0)$ for some $s < t_0$. However, this is clear by inspection: If e_1, e_2, e_3 are the directed edges out of v_0 such that $\{e_1, e_2\}$ is the unique illegal turn in $A(t_0)$, then for $i = 1, 2$ any open edge of $A(s)$ whose image contains e_i must also contain e_3. Since there must be such edges of $A(s)$ for some $s < t_0$, we have that the turns $\{e_1, e_3\}$ and $\{e_2, e_3\}$ are taken, as required. This proves (4) and completes the proof of the lemma. □

We will also require the following lemma which states that along the axis of a principal outer automorphism, bounded length loops are legalized in bounded time. Recall that for a conjugacy class α in F_r and graph $G \in \mathrm{CV}$, $\ell_G(\alpha)$ denotes the length of the immersed representative of α in G.

Lemma 3.2. *Let φ be a principal outer automorphism with lone axis A. For each $l \geq 0$ there is a $D \geq 0$ such that if α is a conjugacy class in F_r such that $\ell_{hA(t_0)}(\alpha) \leq l$ (for some $h \in \mathrm{Out}(F_r)$), then the immersed representative of α in $hA(t)$ is legal for all $t \geq t_0 + D$.*

Proof. By applying the isometry $h \in \mathrm{Out}(F_r)$ of CV, it suffices to prove the lemma for $h = 1$.

There is some $t_1 \in [t_0, t_0 + \ln\lambda]$ such that the folding map $A(t_1) \to A(t_1 + \ln\lambda) = A(t_1)\cdot\varphi$, which we relabel as $f\colon \Lambda \to \Lambda$, is a train track representative of φ mapping vertices to vertices. According to [AKKP18, Proposition 4.11], since φ is principal there are no periodic Nielsen paths in Λ. Hence we may apply [BF94, Proposition 3.1], which states that for any loop β in Λ there is an $N_\beta \geq 0$ such that $[f^{N_\beta}(\beta)]$ (i.e. the tightened image of $f^{N_\beta}(\beta)$ in Λ) is legal. Let

$$N = \max\{N_\beta \colon \ell_{A(t_1)}(\beta) \leq l\}.$$

Then our proof is completed by setting $D = (\ln \lambda)(N+1)$. $\square$

Let (G_t) be a geodesic folding path. For the statement of the next proposition, we say that a nondegenerate turn in G_a is *being folded* (at time $t = a$) if the image of the turn under the folding maps $G_a \to G_b$ is degenerate for *any* $b > a$.

We can now prove our rigidity result concerning folding paths that fellow travel the lone axis A.

Proposition 3.3. *Suppose that $\varphi \in \mathrm{Out}(F_r)$ is a principal outer automorphism with lone axis A. Then there exist constants $\epsilon_0, K_0 \geq 0$ so that if (G_t), for $t \in [t_1, t_2]$, is a greedy geodesic folding path in* CV *that ϵ_0 fellow travels $A' = hA$ (for some $h \in \mathrm{Out}(F_r)$), then the following holds: For any $t \in (t_1 + K_0, t_2)$ and $s \in \mathbb{R}$ such that*

- *G_t is trivalent,*
- *$A'(s)$ is trivalent and in the same open simplex as G_t, and*
- *$\phi_s \colon A'(s) \to G_t$ is a rescaling homeomorphism topologically identifying these graphs,*

we have that a turn in $A'(s)$ is being folded if and only if its image under ϕ_s is being folded in G_t. Hence, ϕ_s preserves the train track structures in the sense that it maps legal turns to legal turns.

We remark that here and throughout, fellow traveling in CV is always meant with respect to the symmetric metric.

Proof. By applying the appropriate isometry $h \in \mathrm{Out}(F_r)$, we note that it suffices to prove the proposition for $A' = A$.

Begin by choosing $\epsilon_0 \leq \log(2)$ so that (G_t) passes through the same sequence of open maximal simplices as A. Also, fix $D \geq 0$, provided by Lemma 3.2, to be such that any loop in $A(t)$ of length no more than 4 is legal in $A(t + D)$.

Let α be a conjugacy class of F_r represented by a legal loop in G_{t_1} such that $\ell_{G_{t_1}}(\alpha) \leq 2$. (Such an α is sometimes called a legal candidate in the literature.) Since $\epsilon_0 \leq \log(2)$, there is a $s_1 \in \mathbb{R}$ such that $d_{\mathrm{sym}}(G_{t_1}, A(s_1)) \leq \epsilon_0 \leq \log(2)$, and so $\ell_{A(s_1)}(\alpha) \leq 4$. By our choice of D in the above paragraph, α is legal in $A(s)$ for all $s \geq s_1 + D$. Moreover, there is a constant $D_2 \geq D$, depending only on the axis A, such that α crosses all legal turns in $A(s)$ for all $s \geq s_1 + D_2$ when $A(s)$ is trivalent. This is because when $A(s)$ is trivalent, all legal turns are taken (Lemma 3.1), and so the difference $D_2 - D$ depends only on the stretch factor of g and the power necessary to make sure that every legal turn is taken by the image of some edge.

Hence, for all trivalent $A(s)$ with $s \geq s_1 + D_2$, α crosses all of the legal turns in $A(s)$ and so α crosses all but the unique illegal turn. If $t \in [t_1, t_2]$ is such that G_t lies in the same open maximal simplex as $A(s)$, then α, which is legal in G_t, crosses all but one turn in G_t. This conclusion holds because $\phi_s \colon A(s) \to G_t$ is a homeomorphism and so maps the immersed representative of α in $A(s)$ to the immersed representative of α in G_t. Hence, the one turn in G_t not taken by α must be the unique illegal turn in G_t. This implies that $\phi_s\colon A(s) \to G_t$ preserves legality, whenever $s \geq s_1 + D_2$ and $A(s)$ and G_t are in the same maximal open simplex.

To complete the proof of the proposition, it suffices to find a $K_0 \geq 0$ such that if $t > K_0 + t_1$, then any $A(s)$ in the same maximal open simplex with G_t necessarily has $s \geq s_1 + D_2$. For this, let $0 < \epsilon$ be the minimum injectivity radius (i.e. length of shortest essential loop) along the periodic line A. Note that if the Lipschitz distance from G_t to a graph in A is less than ϵ_0, then the injectivity radius of G_t is at least $e^{-\epsilon_0}\epsilon$. By compactness, the diameter of the subspace of a simplex consisting of graphs with injectivity radius at least $e^{-\epsilon_0}\epsilon$ is bounded by some constant $\mathfrak{D} \geq 0$. Then setting $K_0 = D_2 + \mathfrak{D} + 2\epsilon_0$ completes the proof by the triangle inequality. □

4. Valencies of branch-points and eventually legalizing folding lines

We begin by stating a convention that we will refer to throughout this section.

Convention 4.1. For the remainder of this section, we assume that $[T] \in \partial\mathrm{CV}$ is given by a *free* F_r-tree T (where $r \geq 3$), that $G_0 \in \mathrm{CV}$, and that $f\colon T_0 \to T$ is an optimal morphism from $T_0 = \widetilde{G}_0$ to T. This data produces the greedy isometric folding path $(\widehat{G}_s)_{s\in J}$ in $\widehat{\mathrm{CV}}$ determined by f starting at $\widehat{G}_0 = G_0$.

Recall from Section 2.4 that the folding path $(\widehat{G}_s)_{s\in J}$ comes together with optimal morphisms $f_s\colon T_s = \widetilde{\widehat{G}_s} \to T$ (where $s \in J$), with "folding maps" $\widehat{g}_{s,s'}\colon \widehat{G}_s \to \widehat{G}_{s'}$ for all $s, s' \in J, s \leq s'$, and their lifts $f_{s,s'}\colon T_s \to T_{s'}$ such that $f_{s'} \circ f_{s,s'} = f_s$. We also have the corresponding geodesic folding path $(G_t)_{t\in J}$ in CV.

Finally, recall that each $\widehat{G}_s$ is given the pullback train track structure $\mathcal{T}_s$ defined by the map f_s; although we note that because the folding path is greedy, the gate structure is unambiguous. By part (2)(a) of Proposition 2.1, the interval J has the form $[0, M)$ for some real number $M > 0$.

We record the following useful general property of our folding paths.

Lemma 4.2. *Let T, $f\colon T_0 \to T$, and $(\widehat{G}_s)_{s\in J}$ be as in Convention 4.1. Let $s \in J$ and let $x \in T_s$ be a vertex with $k \geq 3$ gates with respect to $\mathcal{T}_s$. Then $p = f_s(x) \in T$ is a branch-point with* $\mathrm{val}_T(p) \geq k \geq 3$.

Proof. Let $e_1, \dots, e_k$ be edges of T_s originating at x and representing the k distinct gates at x. Then f_s maps each e_i isometrically to a nondegenerate geodesic segment $f_s(e_i) = [p, p_i]_T$ in T. For $i \neq j$ the edges e_i, e_j are in different gates; therefore the turn $\{e_i, e_j\}$ is legal and the path $e_i^{-1}e_j$ is mapped by f_s injectively to T. This means that for $i = 1, \dots, k$ the segments $[p, p_i]_T$ represent k distinct directions at p in T. Hence $\mathrm{val}_T(p) \geq k \geq 3$, as required. □

Lemma 4.2 motivates the following definition:

Definition 4.3 (Representing branch-points)**.** Let T, $f\colon T_0 \to T$, and $(\widehat{G}_s)_{s\in J}$ be as in Convention 4.1. Let $s \in J$ and let $x \in T_s$ be a vertex with $k \geq 3$ gates with respect to $\mathcal{T}_s$, and let $x_0 \in V\widehat{G}_s$ be the projection of x to $\widehat{G}_s$. Let $p = f_s(x) \in T$ (so that, by Lemma 4.2, p is a branch-point of T of valency $\geq k$).

In this case we say that the branch-point $p \in T$ is *represented* by x, and that the F_r-orbit of p is is *represented* by x_0.

If, moreover, $\mathrm{val}_T(p) = k$, we say that the branch-point $p \in T$ is *faithfully represented* by x, and that the F_r-orbit of p is *faithfully represented* by x_0.

Remark 4.4. Note that if a branch-point $p \in T$ is represented (resp. faithfully represented) by $x \in T_s$ then for each $s' > s$ in J, the branch-point p is also represented (resp. faithfully represented) by $f_{s,s'}(x) \in T_{s'}$.

In general it can happen that in the setting of Lemma 4.2 the point $p = f_s(x) \in T$ has some extra directions not coming from the gates at x in T_s, that is, that $\mathrm{val}_T(p) > k$, so that p is represented but not faithfully represented by x. (For experts: this is exactly what happens in the presence of periodic INPs in train track maps representing some nongeometric fully irreducible $\phi \in \mathrm{Out}(F_r)$.)

Below we define an additional condition satisfied by some "good" folding paths, which will allow us to control and ultimately rule out this kind of behavior. This condition on folding lines is a central point of this paper.

Definition 4.5 (Eventually legalizing folding paths)**.** Let T, $f\colon T_0 \to T$, and $(\widehat{G}_s)_{s\in J}$ be as in Convention 4.1. We say that the folding path $(\widehat{G}_s)_{s\in J}$ is *eventually legalizing* if for any $s \in J$ and any immersed finite path γ in $\widehat{G}_s$, there exists $s' \in J, s' > s$ such that the tightened form $\gamma' = [g_{s,s'}(\gamma)]$ of the image of γ in $\widehat{G}_{s'}$ is legal (with respect to $\mathcal{T}_{s'}$). In this situation we also say that the greedy geodesic folding path $(G_t)_{t\in J'}$ in CV determined by f is *eventually legalizing.*

Note that under the assumptions of Convention 4.1, for every $s \in J$ the subset $f_s(T_s) \subseteq T$ is an F_r-invariant subtree and therefore $f_s(T_s) = T$ since the action of F_r on T is minimal.

Proposition 4.6. *Let T, $f\colon T_0 \to T$, and $(\widehat{G}_s)_{s\in J}$ be as in Convention 4.1. Assume that the greedy isometric folding path $(\widehat{G}_s)_{s\in J}$ is eventually legalizing.*

Then for each branch-point $p \in T$ there exists some $s \in J$ and a vertex $x_0 \in \widehat{G}_s$ such that x_0 faithfully represents the F_r-orbit of p.

Proof. Recall that, by the result of Gaboriau and Levitt, since T is a free F_r-tree, every branch-point of T has finite valency, and there are only finitely many F_r-orbits of branch-points in T (see Section 2.3).

Let $p \in T$ be a branch-point. Thus $3 \leq \mathrm{val}_T(p) = m < \infty$. Let $q_1, \dots, q_m$ be points in T distinct from p such that the directions at p defined by geodesic segments $[p, q_1]_T, \dots, [p, q_m]_T$ represent all m directions at p. In particular, $[p, q_i]_T \cap [p, q_j]_T = \{p\}$ for all $i \neq j$.

Recall that $T_0 = \widetilde{\widehat{G}_0}$ and that $f = f_0 : T_0 \to T$ is onto. Let $u, y_1, \dots, y_m \in T_0$ be such that $f_0(u) = p$ and $f_0(y_i) = q_i$. Denote $\beta_i = [u, y_i]_{T_0}$ and denote by α_i the image of β_i in $\widehat{G}_0$. Thus each α_i is an immersed path in $\widehat{G}_0$ from some point v (the image of u in $\widehat{G}_0$) to some point z_i (the image of y_i in $\widehat{G}_0$). Note that $f_0(\beta_i)$ is a path in T from p to q_i, and so this path passes over $[p,q_i]_T$ but we cannot claim yet that $f_0(\beta_i) = [p,q_i]_T$.

Since our folding path is eventually legalizing, there exists some $s > 0$ in J such that for $i = 1,\dots,m$ the tightened $\widehat{g}_{0,s}$-image τ_i of α_i in $\widehat{G}_s$ is legal. All τ_i have the same initial point v' which is the image of v in $\widehat{G}_s$.

Let $x \in T_s = \widetilde{\widehat{G}_s}$ be a lift of v', and choose a lift ω_i of τ_i to T_s starting at x, for $i = 1,\dots,m$. The fact that τ_i is legal means that $f_s \colon T_s \to T$ is injective on ω_i. Then $f_s(x) = p$ and $f_s(\omega_i) = [p,q_i]_T$ for $i = 1,\dots,m$.

Since we chose $q_1,\dots,q_m$ so that the directions at p in T defined by $[p,q_1]_T,\dots,[p,q_m]_T$ are distinct, the directions defined by $\omega_1,\dots,\omega_m$ at x have to be distinct as well. Otherwise, there would be some $i \neq j$ such that $\omega_i \cap \omega_j$ is nontrivial. But then the image of this overlap $f_s(\omega_i \cap \omega_j)$ would be nontrivial as well, implying that $[p,q_i]_T \cap [p,q_j]_T$ is nontrivial. (Recall that $f_s(\omega_i) = [p,q_i]_T$ and $f_s(\omega_j) = [p,q_j]_T$.) This contradicts our choice of distinct directions at p.

Since $m \geq 3$, this means that x is a vertex of T_s, and hence v' is a vertex of $\widehat{G}_s$, and that the directions at v' represented by initial germs of $\tau_1,\dots,\tau_m$ are in m distinct gates for $\mathcal{T}_s$.

If v' has $k > m$ gates in $\widehat{G}_s$, that would imply that there is another direction at x in T_s which maps by f_s to a direction at p different from the m directions given by $[p,q_1]_T,\dots,[p,q_m]_T$, contradicting the choice of m and of $q_1,\dots,q_m$. Hence v' has exactly m directions in $\widehat{G}_s$. Thus the vertex $x \in T_s$ faithfully represents the branch-point $p \in T$, and the vertex $v' \in \widehat{G}_s$ faithfully represents the F_r-orbit of p, as required. □

We now come to the main result of this section.

Theorem 4.1. *Let $[T] \in \partial\mathrm{CV}$ be a free F_r-tree (where $r \geq 3$), let $T_0 \in \mathrm{CV}$, let $f\colon T_0 \to T$ be an optimal morphism, and let $(\widehat{G}_s)_{s\in J}$ be a greedy isometric folding path in $\widehat{\mathrm{CV}}$ determined by f starting at T_0. Suppose that:*

(1) The folding path $(\widehat{G}_s)_{s\in J}$ is eventually legalizing and
(2) for each $s \in J$ there exists some $s' > s$ in J such that the graph $\widehat{G}_{s'}$ is trivalent.

Then T is trivalent and nongeometric.

Proof. Let $p \in T$ be a branch-point. Then by Proposition 4.6 there exists some $s \in J$ and a vertex $x_0 \in \widehat{G}_s$ such that x_0 faithfully represents the F_r-orbit of p. Thus $\mathrm{val}_T(p) = k \geq 3$, and $\widehat{G}_s$ has exactly k gates at x_0 for $\mathcal{T}_s$. By condition (2), there exists some $s' > s$ in J such that the graph $\widehat{G}_{s'}$ is trivalent. Then, by Remark 4.4, $x_0' = \widehat{g}_{s,s'}(x_0) \in \widehat{G}_{s'}$ is also a vertex with $k \geq 3$ gates that faithfully represents the F_r-orbit of p, and thus $k \leq \deg_{\widehat{G}_{s'}}(x_0')$. Since $\widehat{G}_{s'}$ is trivalent, it follows that $k = 3$. Thus T is trivalent, as required.

We now claim that T is nongeometric. Suppose on the contrary that T is geometric. Then the geometric index of T is equal to $2r - 2$.

Since T is trivalent, and every F_r-orbit trivalent branch-point contributes $3-2=1$ to the geometric index of T, this means that T has exactly $2r-2$ F_r-orbits of branch-points, each of valency 3. Let $p_1, p_2, \ldots, p_{2r-2} \in T$ be representatives of these $2r-2$ F_r-orbits of branch-points in T.

By applying Proposition 4.6, Remark 4.4 and assumption (2), we can find a big enough $s \in J$ such that $\widehat{G}_s$ is trivalent and such that for every $i = 1, \ldots, 2r-2$ there exists a vertex v_i in $\widehat{G}_s$ which faithfully represents the F_r-orbit of p_i and has exactly 3 gates for $\mathcal{T}_s$. The Euler characteristic count for $\widehat{G}_s$ gives us $\sum_v [(\deg(v)/2-1] = r-1$. We also have $\sum_{i=1}^{2r-2}[\deg(v_i)/2-1] = (2r-2)(1/2) = r-1$, which implies that $\widehat{G}_s$ has no other vertices and that $V\widehat{G}_s = \{v_1, \ldots, v_{2r-2}\}$. Since each v_i has degree 3 and has 3 gates in $\widehat{G}_s$, it follows that all non-degenerate turns at v_i are legal for $i = 1, \ldots, 2r-2$, so that all non-degenerate turns in $\widehat{G}_s$ are legal for $\mathcal{T}_s$. This means that $f_s \colon T_s \to T$ is locally injective, and hence an isometry, contradicting the assumption that $[T] \in \partial\mathrm{CV}$. Thus T is nongeometric, as claimed. □

The following lemma characterizes, for an eventually legalizing isometric folding line, how different vertices of $\widehat{G}_s$ can represent branch-points of T belonging to the same F_r-orbit.

Lemma 4.7. *Let T, $f \colon T_0 \to T$, and $(\widehat{G}_s)_{s\in J}$ be as in Convention 4.1. Assume that the greedy isometric folding path $(\widehat{G}_s)_{s\in J}$ is eventually legalizing. Let $s \in J$ and let $x, y \in T_s$ be vertices with ≥ 3 gates which are respectively lifts of vertices $x_0, y_0 \in \widehat{G}_s$. Let $p = f_s(x), q = f_s(y) \in T$ (so that, by Lemma 4.2, p and q are branch-points of T). Then the following are equivalent:*

1. *(1) We have $F_r p = F_r q$.*
2. *(2) There exists some $s' > s$ in J such that $\widehat{g}_{s,s'}(x_0) = \widehat{g}_{s,s'}(y_0)$.*
3. *(3) There exists some $s' > s$ in J and an immersed path γ from x_0 to y_0 in $\widehat{G}_s$ such that the tightened image $[\widehat{g}_{s,s'}(\gamma)]$ of γ in $\widehat{G}_{s'}$ is a trivial path.*

Proof. Note that (3) directly implies (2). And (2) implies (1) as follows. Assume that (2) holds and that $z_0 = \widehat{g}_{s,s'}(x_0) = \widehat{g}_{s,s'}(y_0)$. Recall that we are also given a lift $f_{s',s} \colon T_s \to T_{s'}$ of $\widehat{g}_{s,s'}$ such that $f_s = f_{s'} \circ f_{s',s}$. Then $z_1 = f_{s,s'}(x)$ and $z_2 = f_{s,s'}(y)$ are both lifts of $z_0 = \widehat{g}_{s,s'}(x_0) = \widehat{g}_{s,s'}(y_0)$. We have $p = f_s(x) = f_{s'} \circ f_{s,s'}(x) = f_{s'}(z_1)$ and $q = f_s(y) = f_{s'} \circ f_{s,s'}(y) = f_{s'}(z_2)$. Since both z_1, z_2 are lifts of z_0, it follows that $z_2 = wz_1$ for some $w \in F_r$. Since $p = f_{s'}(z_1)$ and $q = f_{s'}(wz_1)$ and since $f_{s'}$ is F_r-equivariant, we conclude that $q = wp$, and (1) holds.

Finally, suppose that (1) holds and $F_r p = F_r q$. Then there exists $w \in F_r$ such that $q = wp$. Now $f_s(wx) = wf_s(x) = wp = q$. Let γ be the projection to $\widehat{G}_s$ of the geodesic $[y, wx]_{T_s}$. Note that $f_s(y) = f_s(wx) = q$ in T. Since our folding path is eventually legalizing, there exists some $s' > s$ in J such that the tightened path $\gamma' = [\widehat{g}_{s,s'}(\gamma)]$ is legal in $\widehat{G}_{s'}$. If γ' is a nontrivial path, then γ' lifts to a legal immersed path of positive length from $f_{s,s'}(y)$ to $f_{s,s'}(wx)$ in $T_{s'}$ which maps isometrically by $f_{s'}$ to a path of positive length in T from $f_{s'}(f_{s,s'}(y))$ to $f_{s'}(f_{s,s'}(wx))$. This contradicts the fact that $f_{s'}(f_{s,s'}(y)) = f_{s'}(f_{s,s'}(wx)) = q$. Thus γ' is a trivial path in $\widehat{G}_{s'}$. Thus we have proved that (1) implies (3), completing the proof of the lemma. □

In the setting of Convention 4.1, for $s \in J$ let $V'_s \subseteq V\widehat{G}_s$ be the set of all vertices of $\widehat{G}_s$ with ≥ 3 gates for $\mathcal{T}_s$. Define a relation $\sim_s$ on V'_s by setting $v_1 \sim_s v_2$ (for $v_1, v_2 \in V'_s$) if and only if there exists $s' > s, s' \in J$ such that $\widehat{g}_{s,s'}(v_1) = \widehat{g}_{s,s'}(v_2)$ in $\widehat{G}_{s'}$. It is easy to see that $\sim_s$ is an equivalence relation on V'_s. Note that if $v_1 \sim_s v_2$ and v_1 represents the F_r-orbit of a branch-point $p \in T$ then v_2 also represents the F_r-orbit of p, and $\mathrm{val}_T(p) \geq \max\{d_1, d_2\}$ where d_i is the number of gates at v_i in $\widehat{G}_s$ for $i = 1, 2$. However, in this situation if we also have that v_1 faithfully represents the F_r-orbit of $p \in T$, that does not necessarily imply that v_2 faithfully represents the F_r-orbit of $p \in T$ (since it may happen that the number of gates at v_2 is smaller than the number of gates at v_1). For a vertex $v \in V'_s$ we say that v is *maximal for* $\sim_s$ if v has the maximal number of gates among all vertices of V'_s in the $\sim_s$-equivalence class of v.

Corollary 4.8. *Let T, $f\colon T_0 \to T$, and $(\widehat{G}_s)_{s\in J}$ be as in Convention 4.1. Assume that the greedy isometric folding path $(\widehat{G}_s)_{s\in J}$ is eventually legalizing. Let $p_1, \dots, p_m \in T$ be representatives of all the distinct F_r-orbits of branch-points.*

There exists $s_0 \in J$ such that for all $s \geq s_0$ with $s \in J$ the following holds:

(1) There are exactly m distinct $\sim_s$-equivalence classes in V'_s.
(2) Let $v_1, \dots, v_m \in V'_s$ be representatives of all the distinct $\sim_s$-equivalence classes in V'_s, such that for each $i = 1, \dots, m$ the vertex v_i is maximal for $\sim_s$. Then, up to re-ordering of $p_1, \dots, p_m$, for each $i = 1, \dots, m$ the vertex v_i faithfully represents the F_r-orbit of the branch-point p_i of T.

In particular, if k_i is the number of gates at v_i in $\mathcal{T}_s$ then $k_i = \mathrm{val}_T(p_i)$ and

$$\mathrm{ind}_{geom}(T) = \sum_{i=1}^{m} [k_i - 2].$$

Proof. Proposition 4.6 implies that there exists an $s \in J$ such that there are vertices $u_1, \dots, u_m \in V'_s$ where, for each i, we have that u_i faithfully represents the F_r-orbit of p_i. Thus if k_i is the number of gates at u_i in $\mathcal{T}_s$ then $k_i = \mathrm{val}_T(p_i) \geq 3$ for $i = 1, \dots, m$. Since $p_1, \dots, p_m$ are in distinct F_r-orbits, Lemma 4.7 implies that for $i \neq j$ we have $u_i \not\sim_s u_j$. By Lemma 4.2, every vertex $v \in V'_s$ represents the F_r-orbit of some p_i, and therefore, by Lemma 4.7, $v \sim_s u_i$ for some i. Thus there are no other $\sim_s$-equivalence classes in V'_s except the m distinct classes given by $u_1, \dots, u_m$. This means that there are exactly m distinct $\sim_s$-equivalence classes in V'_s, concluding the proof of (1). Moreover, each u_i is maximal in its $\sim_s$-equivalence class, since otherwise there would exist a vertex in V'_s with $> k_i$ gates representing the F_r-orbit of p_i, contradicting the fact that $k_i = \mathrm{val}_T(p_i)$. Thus the conclusion of part (2) in V'_s holds for any maximal elements $v_1, \dots, v_m$ in the $\sim_s$-equivalence classes of $u_1, \dots, u_m$. Remark 4.4 and Lemma 4.7 now imply that the conclusion of part (2) also holds for any $s' > s$ with $s' \in J$. □

Corollary 4.8 provides a precise abstract description of how an eventually legalizing folding path captures the geometric index and the index list for the free F_r-tree $[T] \in \partial\mathrm{CV}$.

5. Random folding rays and principal recurrence

Fix a principal outer automorphism $\varphi \in \mathrm{Out}(F_r)$ with lone axis A in CV.

Definition 5.1 (Recurrent folding rays)**.** A geodesic folding ray (G_t) is *φ-recurrent*, for some principal outer automorphism φ, if there is a $K \geq 0$ such that for any $L \geq 0$, the ray (G_t) has a subsegment that K fellow travels an $\mathrm{Out}(F_r)$-translate of A for length at least L.

We also say that (G_t) is *principally recurrent* if it is φ-recurrent for some principal $\varphi \in \mathrm{Out}(F_r)$.

The main proposition of this section is the following. It is deduced from facts about random walks on groups acting on hyperbolic space (mainly results of Maher–Tiozzo [MT14]) and the bounded geodesic image property for translates of the axis A, a result previously established by the authors [KMPT18].

Proposition 5.2. *Suppose that μ is as in Theorem 2.1 and that φ^{-1} is in the semigroup generated by the support of μ. Let ν be the corresponding hitting measure on* ∂CV *(see Theorem 2.1). Then for ν almost every tree $T \in \partial$*CV *and any geodesic folding ray (G_t) converging to T, we have that (G_t) is φ-recurrent.*

We remind the reader that if φ is principal with axis A in CV, then $\varphi^{-1}A(t) = A(t + \ln \lambda)$. That is, with respect to the left action on CV, φ^{-1} translates A in its folding direction.

Before turning to the proof of Proposition 5.2, we briefly discuss random walks and hyperbolic spaces. The reader can find additional details in [MT14] and a similar setup in [KMPT18]. We assume throughout that μ is a probability measure on G with finite support, although this condition is far stronger than what is needed in this section.

Let X be a δ-hyperbolic space. Given $\kappa > 0$, a quasigeodesic segment $\gamma\colon J \to X$, and a quasigeodesic $\gamma'\colon I \to X$, we say that γ' *crosses* γ *up to distance* κ if there exists an increasing map $\theta\colon J \to I$ such that $d_X(\gamma(t), \gamma'(\theta(t))) \leq \kappa$ for all $t \in J$.

Now suppose we have an isometric action $G \curvearrowright X$ and let γ and γ' be quasigeodesics in X. We say that γ and γ' have an (L, κ)*–oriented match* if there is a subpath $s \subset \gamma$ of diameter at least L and some $h \in G$ such that $h \cdot \gamma'$ crosses s up to distance κ.

Recall that a measure μ on G is nonelementary for the action $G \curvearrowright X$ if the semigroup generated by the support of μ contains 2 loxodromic elements with distinct endpoints on ∂X. Suppose that μ is a nonelementary measure for $G \curvearrowright X$ and that $\varphi \in G$ is a loxodromic in the semigroup generated by the support of μ. In this setting, there is a unique μ-stationary measure ν on ∂X, and ν is the hitting measure for the orbit of the random walk [MT14, Theorem 1.1]. With this setup, we have the following lemma:

Lemma 5.3. *For all $\delta \geq 0$ and all $Q \geq 1$ there is a $\kappa \geq 0$ such that the following holds: For any countable group G acting on a δ-hyperbolic space X, with μ a nonelementary probability measure on G with finite support and hitting measure ν on ∂X, then for ν-almost every $\eta \in \partial X$ and each Q-quasigeodesic ray $\gamma = [x_0, \eta)$ in X with endpoint*

η, the quasigeodesic ray γ has, for each $L \geq 0$, an (L,κ)–oriented match with a Q–quasiaxis α_φ of φ.

Proof. Consider the bi-infinite step space $(G,\mu)^{\mathbb{Z}}$. Let $S\colon (g_n)_{n\in\mathbb{Z}} \mapsto (g_{n+1})_{n\in\mathbb{Z}}$ be the shift map, which acts ergodically on the step space. Let $w\colon (g_n)_{n\in\mathbb{Z}} \mapsto (w_n)_{n\in\mathbb{Z}}$ be the map from the step space to the path space $(G^{\mathbb{Z}}, \mathbb{P})$, where

$$w_n = \begin{cases} g_1 g_2 \dots g_n & \text{for } n > 0 \\ g_0^{-1} g_1^{-1} \dots g_{-n+1}^{-1} & \text{for } n \leq 0, \end{cases}$$

and $\mathbb{P}$ is the push forward of the product measure $\mu^{\mathbb{Z}}$ by w. By [MT14], almost every sample path converges in both the forward and backward directions, giving rise to a map $\partial = \partial_+ \times \partial_- \colon (G^{\mathbb{Z}}, \mathbb{P}) \to \partial X \times \partial X$, defined on a full measure subset of the path space. In particular, this means that the shift map S acts ergodically on $(G^{\mathbb{Z}}, \mathbb{P})$, where $S^k(w_n)_{n\in\mathbb{Z}} = (w_k^{-1} w_n)$. Furthermore, $\nu \times \check{\nu}$, the product of the hitting measure with the reflected hitting measure, is the push forward of the path space measure $\mathbb{P}$ under ∂.

Given an oriented Q-quasiaxis α_φ, we shall write α_φ^+ and α_φ^- for its forward and backward limit points in ∂X respectively. We shall write $\alpha_\varphi(0)$ for a nearest point on α_φ to the basepoint x_0 in X. Given constants $\delta \geq 0$ and $Q \geq 0$, there is a constant $\kappa \geq 0$, such that for any Q-quasigeodesic α_φ in a δ-hyperbolic space, and any constant $L \geq 0$, there are open sets A and B in ∂X, with $\alpha_\varphi^- \in A$ and $\alpha_\varphi^+ \in B$ such that any bi-infinite Q-quasigeodesic γ, with one endpoint in A and the other in B, κ-fellow travels with a subquasigeodsic of α_φ of length at least L, centered at $\alpha_\varphi(0)$. Furthermore, the distance between $\alpha_\varphi(0)$ and the closest point on γ to the basepoint x_0 is bounded in terms of δ and Q.

We shall write γ_ω to denote a bi-infinite Q-quasigodesic connecting the forward and backward limit points of $(w_n)_{n\in\mathbb{Z}}$. If $S^k(w_n)_{n\in\mathbb{Z}}$ lies in $\partial^{-1}(A\times B)$, then there is a subsegment of γ_ω of length L, centered at the nearest point projection of w_k to γ_ω, which fellow travels with $w_k\alpha_\varphi$. As φ lies in the semigroup generated by the support of μ, by [MT14, Proposition 5.4], $\nu \times \check{\nu}(A \times B) = \nu(A)\check{\nu}(B)$ is strictly positive. In particular, $\partial^{-1}(A \times B)$ is positive. Therefore, by Birkhoff's pointwise ergodic theorem, the proportion of integers $1 \leq k \leq N$ such that $S^k(w_n)_{n\in Z}$ lies in $\partial^{-1}(A \times B)$ converges to $\nu(A)\check{\nu}(B)$ as $N \to \infty$. In particular, there is a sequence of integers $k_i \to \infty$ such that $S^{k_i}(w_n)_{n\in\mathbb{Z}}$ lies in $\partial^{-1}(A\times B)$, and as $(w_n)_{n\in\mathbb{Z}}$ converges to $\partial_+(w_n)_{n\in\mathbb{Z}}$, this means that there are infinitely many disjoint subintervals of γ_ω of length L which κ-fellow travel with a translate of α_φ. The same property now follows for Q-quasigeodesic rays starting at x_0 and converging to $\partial_+(w_n)_{n\in\mathbb{Z}}$, as every such ray has an infinite terminal subray which fellow travels with γ_ω.

So we have shown that for some $\kappa \geq 0$ and any $L \geq 0$, the set of $\eta \in \partial X$ for which any Q-quasigeodesic ray $\gamma = [x_0, \eta)$ has an (L,κ)-oriented match with α_φ has ν measure 1. Intersecting these sets over all $L \in \mathbb{Z}_+$, we see that the set of $\eta \in \partial X$ such that every Q-quasigeodesic ray $\gamma = [x_0,\eta)$ has an (L,κ)–oriented match with α_φ *for every* $L \geq 0$ also has ν measure 1. This completes the proof. □

Now Proposition 5.2 follows from Lemma 5.3 and the bounded geodesic image property for translates of A.

Proof of Proposition 5.2. Recall that for ν-a.e. tree $T \in \partial\mathrm{CV}$, we have that T is free, arational, and uniquely ergodic (Theorem 2.1 and Corollary 2.3.) Hence, by Proposition 2.1, there exists a geodesic folding ray (G_t) converging to T.

The π-image of any geodesic folding path in the free factor complex $\mathcal{FF}$ is a Q-unparameterized quasigeodesic, for Q depending only on the rank of F_r [BF14, Corollary 6.5]. Since φ acts as a loxodromic isometry on $\mathcal{FF}$, at the expense of increasing Q, we may assume that the image $\pi(A)$ of the axis A is a Q-quasiaxis for φ in $\mathcal{FF}$. So applying Lemma 5.3 to the situation at hand, gives that almost surely the quasiray $\pi((G_t))$ has an (L, κ)-oriented match with $\pi(A)$ for every $L \geq 0$.

Unpacking this statement, we see that for any $L \geq 0$, there is an $h \in \mathrm{Out}(F_r)$ such that $\pi((G_t))$ has a segment of diameter at least L in $\mathcal{FF}$ that crosses $\pi(hA)$ up to distance κ. Since the map $\pi\colon \mathrm{CV} \to \mathcal{FF}$ is coarsely Lipschitz [BF14, Corollary 3.5], it suffices to show that fellow traveling of $\pi((G_t))$ and $\pi(hA)$ in $\mathcal{FF}$ can be lifted to uniform fellow traveling of (G_t) and hA in CV. This follows from the bounded geodesic image property established in [KMPT18, Theorem 7.8] and the rest of the argument is similar to the one given for [KMPT18, Theorem A].

In some detail, if $\pi((G_t))$ and $\pi(hA)$ fellow travel for L sufficiently large, then the nearest point projection in $\mathcal{FF}$ of the path $\pi((G_t))$ to $\pi(hA)$ is roughly diameter L, depending only on Q and the hyperbolicity constant of $\mathcal{FF}$. In terms of Outer space, this means that the projection of (G_t) to the greedy folding axis hA using the Bestvina–Feighn (see [BF14]) projection $\mathrm{Pr}_{hA}\colon \mathrm{CV} \to \mathrm{hA}$ is no less than cL, for some $c \geq 0$ depending only on the rank of F_r. This follows from the fact, established in [DT18, Lemma 4.2], that $\pi \circ \mathrm{Pr}_{hA}$ is coarsely equal to $\mathbf{n} \circ \pi$, where $\mathbf{n}\colon \mathcal{FF} \to \pi(hA)$ is the nearest point projection. Corollary 7.9 of [KMPT18] then implies that the path (G_t) contains a subsegment that K fellow travels a subsegment of hA for distance $cL - c_1$, for some constants $c_1, K \geq 0$ that depend only on the principal outer automorphism φ. Since this was true for any $L \geq 0$, we have that (G_t) is φ-recurrent and the proof is complete. $\square$

6. Principally recurrent folding lines are eventually legalizing

In this section, we fix a principal outer automorphism $\varphi \in \mathrm{Out}(F_r)$ and denote by A its lone folding axis in CV. Our goal is to show that principally recurrent folding paths are all eventually legalizing. This is achieved in Proposition 6.2.

Our first lemma is proven in the same manner as Lemma 5.9 of [KMPT18]. It basically states that in the case of interest, if folding paths fellow travel for a long enough time, then they get arbitrarily close to one another.

Lemma 6.1. *If the greedy geodesic folding ray (G_t) is φ-recurrent, then for any $\epsilon > 0$ and any $L \geq 0$, the ray (G_t) has a subsegment that ϵ fellow travels an* $\mathrm{Out}(F_r)$*–translate of A for length at least L.*

Proof. Suppose this were not the case. Then, using the periodicity of A and φ-recurrence of (G_t), we could find a sequence of $h_i \in \mathrm{Out}(F_r)$ and an $\epsilon > 0$ so that the rays $h_i(G_t)$ K fellow travel A about the point $A(t_0)$ for length L_i, but the symmetric

ϵ–ball about $A(t_0)$ does not meet any of the $h_i(G_t)$. Here, $K \geq 0$ is a fixed constant, t_0 is a fixed time, and $L_i \to \infty$ as $i \to \infty$.

Then, just as in the proof of Lemma 5.9 of [KMPT18], the sequence $h_i(G_t)$ has a subsequence that converges to a greedy folding line B which has bounded distance from A (see also [BR15, Lemma 6.11]). In particular, B has the same limit points in ∂CV as A (as in Lemma 3.1.1), but does not contain the point $A(t_0)$. This is to say that B is a folding line from the repelling tree to the attracting tree of φ that is distinct from A, contradicting the fact that φ is a lone axis outer automorphism. □

The main result of this section is the following proposition.

Proposition 6.2. *Suppose that the greedy geodesic folding ray (G_t) in* CV *is φ-recurrent. Then (G_t) is eventually legalizing.*

Proof. Let γ_0 be an immersed path in G_0 and let γ_t denote its image in G_t (via the fold maps) after tightening. In general, if p is any path in G, its tightening is denoted $[p]$. Our goal is to show that γ_t is legal in G_t for sufficiently large t.

Let N be the number of illegal turns in γ_0. Note that the number of illegal turns N_t in γ_t is nonincreasing in t and so $N_t \leq N$. We begin by choosing $\mathfrak{t}_0 \geq 0$ sufficiently large so that for all $t \geq \mathfrak{t}_0$,

- $N_t = N_{\mathfrak{t}_0}$, i.e. the number of illegal turns has stabilized.

Hence, for all $t \geq \mathfrak{t}_0$ we have the decomposition

$$\gamma_t = \gamma_t^0 \cdot \ldots \cdot \gamma_t^{N_{\mathfrak{t}_0}}, \tag{1}$$

where the breakpoints happen exactly at the illegal turns of γ_t. In the language of Section 5 of [BF14], γ_t has all *surviving illegal turns* for the folding ray, in the sense that no illegal turns of γ_t become legal or collide with one another while folding. Although it is not strictly needed for what follows, this observation makes it clear how the decomposition of $\gamma_{t'}$ is obtained from the decomposition of γ_t for $\mathfrak{t}_0 \leq t \leq t'$: just consider the image of γ_t^i under the folding map to $G_{t'}$ and remove initial and terminal portions of the image that cancel with portions of its neighbors. Since the number of illegal turns in γ_t does not decease for $t \geq \mathfrak{t}_0$, these images are never canceled away.

Returning to the argument, by Corollary 4.8 of [BF14], for $s \geq t$ any legal segment σ_t inside of γ_t of length $L_t \geq 2$ gives rise to a legal segment σ_s inside of γ_s of length $L_s \geq 2 + (L_t - 2)e^{s-t}$. (This conclusion follows from the so-called derivative formula of Bestvina–Feighn, [BF14, Lemma 4.4].) Hence, if at any time γ_t^i has length at least 3, then it grows exponentially thereafter. So at the expense of making $\mathfrak{t}_0$ larger, we may additionally assume that for each $0 \leq i \leq N_{\mathfrak{t}_0}$ either:

- $\gamma_{\mathfrak{t}_0}^i$ has length at least 8 (and hence has length ≥ 8 for all $t \geq \mathfrak{t}_0$), or
- γ_t^i has length at most 2 for all $t \geq \mathfrak{t}_0$.

We call the γ_t^is of length greater than 8 *large* and the rest are called *small.*

Note that if $N_{\mathfrak{t}_0} = 0$, then we are done. So assume that $N_{\mathfrak{t}_0} > 0$.

Now for any $s \geq \mathfrak{t}_0$ we use (1) to construct another decomposition of γ_s,

$$\gamma_s = r_s^1 \cdot r_s^2 \cdot \ldots \cdot r_s^k \tag{2}$$

for $k \leq N_t$ defined as follows: for each large γ_s^i there are two breakpoints of the decomposition (2) at vertices along γ_s^i obtained by starting at the endpoints γ_s^i, moving inward (along γ_s^i) for length 2 and choosing the next vertices of γ_s^i (while continuing to move along γ_s^i). Since the length of γ_s^i is at least 8 and every edge has length less than 1, this process chooses two vertex breakpoints per large γ_s^i, and results in a decomposition of γ_s in which each term begins and ends with (possibly overlapping) legal segments of length at least 2.

The decomposition of γ_s given in (2) is a *splitting* in the sense that if we denote the folding maps by $g_{s,t}\colon G_s \to G_t$, we have for $\mathfrak{t}_0 \leq s < t$

$$\gamma_t = [g_{s,t}(r_s^1)] \cdot \ldots \cdot [g_{s,t}(r_s^k)].$$

This again follows from the formulation of the derivative formula stated above since legal segments of length at least 2 are not completely cancelled under folding. (We warn the reader that we are not claiming that the above splitting of γ_t is the same as the one appearing in (2) for $s = t$.)

Note that (for each $s \geq \mathfrak{t}_0$) the r_s^j's alternate between legal segments (of length at least 2) and *clusters* of segments of length no more than 3 joined by illegal turns. The total length of each illegal cluster is no more than $3N_{\mathfrak{t}_0} \leq 3N_0$. Moreover, if r_s^j is an illegal cluster of γ_s, then for any $t > s$, r_t^j is an illegal cluster of γ_t and r_t^j is a subpath of $[g_{s,t}(r_s^j)]$ whose complementary pieces are legal initial/terminal subpaths of $[g_{s,t}(r_s^j)]$. This fact follows directly from our construction.

Since $N_{\mathfrak{t}_0} > 0$ and all illegal turns of γ_s are contained in illegal clusters, there exists a $1 \leq j \leq k$ such that r_s^j is an illegal cluster for all $t \geq \mathfrak{t}_0$. We set $r_s = r_s^j$ and henceforth work only with this illegal cluster. We will show that for some $\mathfrak{t}_0 < s < t$, the immersed path $[g_{s,t}(r_s)]$ is *legal* in G_t. Since this is a subpath of γ_t, this shows that $N_t < N_{\mathfrak{t}_0}$; a contradiction that will complete the proof.

Now apply Lemma 3.2 with $l = 2(3N_0 + 5)$ to obtained a $D \geq 0$ so that for any $t \in \mathbb{R}$ and $h \in \mathrm{Out}(F_r)$, any loop in $hA(t)$ of length at most $2(3N_0 + 5)$ becomes legal in $hA(t + D)$, after folding and tightening. Also fix $\epsilon \leq \min\{\epsilon_0, \log(2)\}$ and $L \geq K_0 + D + 2\epsilon + 2$, where ϵ_0 and K_0 are as in Proposition 3.3. As (G_t) is φ-recurrent, Lemma 6.1 implies that for this $\epsilon, L \geq 0$, there is a interval (after time $\mathfrak{t}_0$) on which (G_t) ϵ fellow travels $hA(t)$ (for some $h \in \mathrm{Out}(F_r)$) for length L. For ease of notation, set $A' = hA$.

Hence, we have obtained a subinterval $[\mathfrak{t}_1, \mathfrak{t}_1 + L]$ ($\mathfrak{t}_1 \geq \mathfrak{t}_0$) such that (G_t) ϵ fellow travels A' for $t \in [\mathfrak{t}_1, \mathfrak{t}_1 + L]$. Applying Proposition 3.3, we get a subinterval $[\mathfrak{t}_1 + K_0, \mathfrak{t}_1 + L]$ of length at least $D + 2\epsilon + 2$ with the property that for any $t \in (\mathfrak{t}_1 + K_0, \mathfrak{t}_1 + L)$ and $s \in \mathbb{R}$ such that

(a) G_t is trivalent,
(b) $A'(s)$ is trivalent and in the same open simplex as G_t, and
(c) $\phi_s\colon A'(s) \to G_t$ is a homeomorphism topologically identifying these graphs,

we have that ϕ_s preserves the train track structures in the sense that it maps legal turns to legal turns.

Finally, choose $a, b \in (\mathfrak{t}_1 + K_0, \mathfrak{t}_1 + L)$ such that:

(i) a is within distance 1 of $\mathfrak{t}_1 + K_0$ and b is within distance 1 of $\mathfrak{t}_1 + L$ (and so $b - a > D + 2\epsilon$),
(ii) G_a, G_b are trivalent, and
(iii) there are $c, d \in \mathbb{R}$ such that $A'(c), A'(d)$ are trivalent, in the same maximal simplex as G_a, G_b (respectively), and the symmetric distance between $A'(c)$ and G_a (and $A'(d)$ and G_b) is no more than ϵ.

Let $\phi_c \colon A'(c) \to G_a$ and $\phi_d \colon A'(d) \to G_b$ be the homeomorphisms preserving the associated train track structures. Since $A'(c)$ has exactly one illegal turn (Lemma 3.1), the same is true for G_a.

Recall that the illegal cluster r_a has length no more than $3N_0$ in G_a. Since there is only one illegal turn of G_a we can easily 'legally' extend r_a to a immersed loop α_a. By this we mean that α_a is an immersed loop containing r_a so that the rest of α_a (call it p_a) is a legal arc of length at least 2 which meets the endpoints of r_a at legal turns. It is also easy to see that can be done in such a way that α_a has length no more than 5 plus the length of r_a.

Let α be the conjugacy class of F_r represented by α_a in G_a and let α_t denote the immersed representative of α in G_t for $t \geq a$. Hence, $\ell_{G_a}(\alpha) \leq 3N_0 + 5$.

We claim that for all $t > a$, $[f_{a,t}(r_a)]$ is a subpath of α_t in G_t. This conclusion is an immediate consequence of the fact that $[f_{a,t}(\alpha_a)] = \alpha_t$ and the fact that

$$\alpha_a = r_a \cdot p_a,$$

is a splitting of α_a (as a loop). This last fact again follows from our construction and the formulation of the Bestvina–Feighn derivative formula used above.

We are now ready to complete the proof of Proposition 6.2. First, since the symmetric distances $d_{\mathrm{sym}}(G_a, A'(c))$ and $d_{\mathrm{sym}}(G_b, A'(d))$ are each less than ϵ and $b - a \geq D + 2\epsilon$, we have (using that (G_t) and A' are directed geodesic) that $d - c \geq D$. Moreover, again using this distance estimate and the fact that $\epsilon \leq \log(2)$,

$$\ell_{A'(c)}(\alpha) \leq 2\ell_{G_a}(\alpha) \leq 2(3N_0 + 5).$$

Our choice of D, then gives that the immersed representative of α in $A'(c + D)$ is legal, and hence it is legal in $A'(d)$. But since the homeomorphism $\phi_d \colon A(d) \to G_b$ maps the immersed representative of α in $A(d)$ to the immersed representative of α in G_b and preserves legality, the immersed representative of α in G_b is legal. This is all to say that α_b is a legal loop in G_b. Since α_b contains the path $[f_{a,b}(r_a)]$, this path too is legal in G_b. But this is exactly the contradiction we sought, and so the proof of Proposition 6.2 is complete. □

7. Proof of the main result

Recall that a probability measure μ on $\mathrm{Out}(F_r)$ is called *nonelementary* if the subsemigroup $\langle \mathrm{Supp}(\mu) \rangle_+$ of $\mathrm{Out}(F_r)$ generated by the support $\mathrm{Supp}(\mu)$ of μ contains two independent fully irreducible elements (that is, two fully irreducible elements $\psi_1, \psi_2 \in \mathrm{Out}(F_r)$ such that the subgroup $\langle \psi_1, \psi_2 \rangle$ is not virtually cyclic).

We can now prove the main result of this paper (c.f. Theorem 1.1 in the introduction):

Theorem 7.1. *Suppose that $r \geq 3$ and let μ be a nonelementary probability measure on $\mathrm{Out}(F_r)$ with finite support such that $\varphi^{-1} \in \langle \mathrm{Supp}(\mu) \rangle_+$ for some principal fully irreducible $\varphi \in \mathrm{Out}(F_r)$. Let ν be the hitting measure on $\partial\mathrm{CV}$ for the random walk $(\mathrm{Out}(F_r), \mu)$ starting at some $y_0 \in \mathrm{CV}$.*

Then for ν-a.e. $[T] \in \partial\mathrm{CV}$, the tree T is trivalent and nongeometric.

Proof. By Corollary 2.3 and Theorem 2.1, for ν-a.e. $[T] \in \partial\mathrm{CV}$, the tree T is F_r-free and uniquely ergodic.

By Proposition 2.1, there exists a (greedy) geodesic folding ray (G_t) in CV such that $\lim_{t\to\infty} G_t = [T]$ in $\overline{\mathrm{CV}}$. Proposition 5.2 now implies that the ray (G_t) is φ-recurrent. Hence, by Proposition 6.2, the ray (G_t) is eventually legalizing. Therefore, by Theorem 4.1, the tree T is trivalent and nongeometric. □

Corollary 7.1. *Suppose that $r \geq 3$ and let μ be a nonelementary probability measure on $\mathrm{Out}(F_r)$ with finite support such that $\langle \mathrm{Supp}(\mu) \rangle_+$ contains a subgroup of finite index in $\mathrm{Out}(F_r)$. Let ν be the hitting measure on $\partial\mathrm{CV}$ for the random walk $(\mathrm{Out}(F_r), \mu)$ starting at some $y_0 \in \mathrm{CV}$.*

Then for ν-a.e. $[T] \in \partial\mathrm{CV}$, the tree T is trivalent and nongeometric.

Proof. Let $H \leq \mathrm{Out}(F_r)$ be a subgroup of finite index such that $H \subseteq \langle \mathrm{Supp}(\mu) \rangle_+$. By [AKKP18, Example 6.1], there exists a principal fully irreducible $\varphi \in \mathrm{Out}(F_r)$. Then for some $m \geq 1$ we have $\varphi^m \in H$ and therefore $\varphi^{-m} \in \langle \mathrm{Supp}(\mu) \rangle_+$. Hence, by Theorem 7.1 above, the statement of the corollary follows. □

References

[AK11] Y. Algom-Kfir. Strongly contracting geodesics in outer space. *Geom. Topol*, 15(4):2181–2233, 2011.

[AKKP18] Y. Algom-Kfir, I. Kapovich, and C. Pfaff. Stable Strata of Geodesics in Outer Space. *International Mathematics Research Notices*, 2018(00):pp. 1–30, 2018.

[BF94] M. Bestvina and M. Feighn. Outer limits. *preprint*, pages 1–19, 1994.

[BF14] M. Bestvina and M. Feighn. Hyperbolicity of the complex of free factors. *Advances in Mathematics*, 256:104–155, 2014.

[BFH97] M. Bestvina, M. Feighn, and M. Handel. Laminations, trees, and irreducible automorphisms of free groups. *Geometric and Functional Analysis*, 7(2):215–244, 1997.

[BHW16] M. Bestvina, C. Horbez, and R. Wade. On the topological dimension of the Gromov boundaries of some hyperbolic $Out(F_N)$-graphs. *arXiv preprint arXiv:1610.02115*, 2016.

[BR15] M. Bestvina and P. Reynolds. The boundary of the complex of free factors. *Duke Mathematical Journal*, 164(11):2213–2251, 2015.

[CH12] T. Coulbois and A. Hilion. Botany of irreducible automorphisms of free groups. *Pacific Journal of Mathematics*, 256(2), 2012.

[CHL08a] T. Coulbois, A. Hilion, and M. Lustig. R-trees and laminations for free groups I: algebraic laminations. *Journal of the London Mathematical Society*, 78(3):723–736, 2008.

[CHL08b] T. Coulbois, A. Hilion, and M. Lustig. R-trees and laminations for free groups II: the dual lamination of an R-tree. *Journal of the London Mathematical Society*, 2008.

[CL95] M. Cohen and M. Lustig. Very small group actions on r-trees and dehn twist automorphisms. *Topology*, 34(3):575–617, 1995.

[CV86] M. Culler and K. Vogtmann. Moduli of graphs and automorphisms of free groups. *Inventiones mathematicae*, 84(1):91–119, 1986.

[DH17] F. Dahmani and C. Horbez. Spectral Theorems for Random Walks on Mapping Class Groups and $Out(F_n)$. *International Mathematics Research Notices*, page rnw306, 2017.

[DT17] S. Dowdall and S. J. Taylor. The co-surface graph and the geometry of hyperbolic free group extensions. *Journal of Topology*, 10:447–482, 2017.

[DT18] S. Dowdall and S. J. Taylor. Hyperbolic extensions of free groups. *Geom. Topol.*, 22(1):517–570, 2018.

[FM11] S. Francaviglia and A. Martino. Metric properties of outer space. *Publicacions Matemàtiques*, 55(2):433–473, 2011.

[GL95] D. Gaboriau and G. Levitt. The rank of actions on R-trees. In *Annales scientifiques de l'Ecole normale supérieure*, volume 28, pages 549–570. Société mathématique de France, 1995.

[GM17] V. Gadre and J. Maher. Recurrence of quadratic differentials for harmonic measure. *arXiv preprint arXiv:1706.01926*, 2017.

[Ham12] U. Hamenstädt. The boundary of the free factor graph. *arXiv preprint arXiv:1211.1630*, 2012.

[HM11] M. Handel and L. Mosher. *Axes in outer space.* Number 1004. Amer Mathematical Society, 2011.

[Hor16] C. Horbez. The poisson boundary of $\mathrm{Out}(F_n)$. *Duke Mathematical Journal*, 165(2):341–369, 2016.

[Hor17] C. Horbez. Central limit theorems for mapping class groups and $\mathrm{Out}(F_N)$. *Geometry & Topology*, 22(1):105–156, 2017.

[KMPT18] I. Kapovich, J. Maher, C. Pfaff, and S. J. Taylor. Random outer automorphisms of free groups: Attracting trees and their singularity structures. *arXiv preprint arXiv:1805.12382*, 2018.

[MP16] L. Mosher and C. Pfaff. Lone axes in outer space. *Algebr. Geom. Topol.*, 16(6):3385–3418, 2016.

[MT14] J. Maher and G. Tiozzo. Random walks on weakly hyperbolic groups. *Journal fur die reine und angewandte Mathematik (Crelles Journal)*, 2014.

[NPR14] H. Namazi, A. Pettet, and P. Reynolds. Ergodic decompositions for folding and unfolding paths in outer space. *arXiv preprint arXiv:1410.8870*, 2014.

[Pau89] F. Paulin. The Gromov topology on **R**-trees. *Topology Appl.*, 32(3):197–221, 1989.

[Rey12] P. Reynolds. Reducing systems for very small trees. *arXiv preprint arXiv:1211.3378*, 2012.

[TT16] S. J. Taylor and G. Tiozzo. Random extensions of free groups and surface groups are hyperbolic. *Int. Math. Res. Not.*, 2016(1):294–310, 2016.

Department of Mathematics and Statistics, Hunter College of CUNY
695 Park Ave, New York, NY 10065
http://math.hunter.cuny.edu/ilyakapo/,
Email address: ik535@hunter.cuny.edu

CUNY College of Staten Island and CUNY Graduate Center
2800 Victory Boulevard, Staten Island, NY 10314
http://www.math.csi.cuny.edu/~maher/,
Email address: joseph.maher@csi.cuny.edu

Department of Math & Stats, Queen's University
Jeffery Hall, 48 University Ave., Kingston, ON Canada, K7L 3N6
https://mast.queensu.ca/~cpfaff/,
Email address: c.pfaff@queensu.ca

Department of Mathematics, Temple University
1805 Broad St, Philadelphia, PA 19122
https://math.temple.edu/~samuel.taylor/,
Email address: samuel.taylor@temple.edu